

Linear Convergence of SDCA in Statistical Estimation

Chao Qu¹ and Huan Xu²

¹Department of Mechanical Engineering, National University of Singapore

²H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology

Abstract

In this paper, we consider stochastic dual coordinate (SDCA) *without* strongly convex assumption or convex assumption. We show that SDCA converges linearly under mild conditions termed restricted strong convexity. This covers a wide array of popular statistical models including Lasso, group Lasso, and logistic regression with ℓ_1 regularization, corrected Lasso and linear regression with SCAD regularizer. This significantly improves previous convergence results on SDCA for problems that are not strongly convex. As a by product, we derive a dual free form of SDCA that can handle general regularization term, which is of interest by itself.

1 Introduction

First order methods have again become a central focus of research in optimization and particularly in machine learning in recent years, thanks to its ability to address very large scale empirical risk minimization problems that are ubiquitous in machine learning, a task that is often challenging for other algorithms such interior point methods. The randomized (dual) coordinate version of the first order method samples one data point and updates the objective function at each time step, which avoids the computations of the full gradient and pushes the speed to a higher level. Related methods have been implemented in various software packages Vedaldi and Fulkerson [2008]. In particular, the randomized dual coordinate method considers the following problem.

$$\begin{aligned} \min_{w \in \Omega} F(w) &:= \frac{1}{n} \sum_{i=1}^n f_i(w) + \lambda g(w) \\ &= f(w) + \lambda g(w), \end{aligned} \tag{1}$$

where $f_i(w)$ is a convex loss function of each sample and $g(w)$ is the regularization, Ω is a convex compact set. Instead of directly solving the primal problem it look at the dual problem

$$D(\alpha) = \frac{1}{n} \sum_{i=1}^n -\psi_i^*(-\alpha_i) - \lambda g^*\left(\frac{1}{\lambda n} \sum_{i=1}^n X_i \alpha_i\right),$$

where it assumes the loss function $f_i(w)$ has the form $\psi_i(X_i^T w)$. Shalev-Shwartz and Zhang [2013] consider this dual form and proved linear convergence of the stochastic dual coordinate ascent method (SDCA) when $g(w) = \|w\|_2^2$. They further extend the result into

a general form, which allows the regularizer $g(w)$ to be a general strongly convex function Shalev-Shwartz and Zhang [2014]. Qu and Richtárik [2015] propose the AdaSDCA, an adaptive variant of SDCA, to allow the method to adaptively change the probability distribution over the dual variables through the iterative process. Their experiment results outperform the non-adaptive methods. Dünner et al. [2016] consider a primal-dual frame work and extend SDCA to the non-strongly convex case with sublinear rate. Allen-Zhu and Yuan [2015a] further improve the convergece speed using a novel non-uniform sampling that selects each coordinate with a probability proportional to the square root of the smoothness parameter. Other acceleration techniques Qu and Richtárik [2014], Qu et al. [2015], Nesterov [2012], as well as mini-batch and distributed variants on coordinate method Liu and Wright [2015], Zhao et al. [2014], Jaggi et al. [2014], Mahajan et al. [2014] have been studied in literature. See Wright [2015] for a review on the coordinate method.

Our goal is to investigate how to extend SDCA to non-convex statistical problems. Non-convex optimization problems attract fast growing attentions due to the rise of numerous applications, notably non-convex M estimators (e.g., SCAD Fan and Li [2001], MCP Zhang [2010]), deep learning Goodfellow et al. [2016] and robust regression [e.g., correct Lasso in Loh and Wainwright, 2011]. Shalev-Shwartz [2016] proposes a dual free version of SDCA and proves the linear convergence of it, which addresses the case that each individual loss function $f_i(w)$ may be non-convex, but their sum $f(w)$ is strongly convex. This extends the applicability of SDCA. From a technical perspective, due to non-convexity of $f_i(w)$, the paper avoids explicitly using its dual (hence the name “dual free”), by introducing pseudo-dual variables.

In this paper, we consider using SDCA to solve M-estimators that are not strongly convex, or even non-convex. We show that under the *restricted strong convexity* condition, SDCA converges linearly. This setup includes well known formulations such as Lasso, Group Lasso, Logistic regression with ℓ_1 regularization, linear regression with SCAD regularizer Fan and Li [2001], Corrected Lasso Loh and Wainwright [2011], to name a few. This significantly improves upon existing theoretic results Shalev-Shwartz and Zhang [2013, 2014], Shalev-Shwartz [2016] which only established sublinear convergence on convex objectives.

To this end, we first adapt SDCA Shalev-Shwartz and Zhang [2014] into a generalized dual free form in Algorithm 1. This is because to apply SDCA in our setup, we need to introduce a non-convex term and thus demand a dual free analysis. We remark that the theory about dual free SDCA established in Shalev-Shwartz [2016] *does not* apply to our setup for the following reasons:

1. In Shalev-Shwartz [2016], $F(w)$ needs to be strongly convex, thus does not apply to Lasso or non-convex M-estimators.
2. Shalev-Shwartz [2016] only studies the special case where $g(w) = \frac{1}{2}\|w\|_2^2$, while the M-estimators we consider include non-smooth regularization such as ℓ_1 or $\ell_{1,2}$ norm.

We illustrate the update rule of Algorithm 1 using an example. While our main focus is for non-strongly convex problems, we start with a strongly convex example to obtain some intuitions. Suppose $f_i(w) = \frac{1}{2}(y_i - w^T x_i)^2$ and $g(w) = \frac{1}{2}\|w\|_2^2 + \|w\|_1$. It is easy to see $\nabla f_i(w) = x_i(w^T x_i - y_i)$ and $w_t = \arg \min_{w \in \Omega} \frac{1}{2}\|w - v_t\|_2^2 + \|w\|_1$.

It is then clear that to apply SDCA $g(w)$ needs to be strongly convex, or otherwise the proximal step $w_t = \arg \max_{w \in \Omega} w^T v^t - \|w\|_1$ becomes ill-defined: w_t may be infinity (if Ω is unbounded) or not unique. This observation motivates the following preprocessing

Algorithm 1 Generalized Dual free SDCA

Input: Step size η , number of iterations T , and smoothness parameters L_1, \dots, L_n
Initialization: $v^0 = \frac{1}{\lambda n} \sum_{i=1}^n a_i^0$ for some $a^0 = (a_1^0, \dots, a_n^0)$, $q_i = \frac{L_i + \bar{L}}{2n\bar{L}}$, where $\bar{L} = \frac{1}{n} \sum_{i=1}^n L_i$, probability distribution $Q = (q_1, q_2, \dots, q_n)$ on $\{1, \dots, n\}$.
for $t = 1, \dots, T$ **do**
 pick i randomly according to Q , $\eta_i = \frac{\eta}{q_i n}$;
 update: $a_i^t = a_i^{t-1} - \eta_i \lambda n (\nabla f_i(w^{t-1}) + a_i^{t-1})$;
 update: $v^t = v^{t-1} - \eta_i (\nabla f_i(w^{t-1}) + a_i^{t-1})$;
 update: $w^t = \arg \max_{w \in \Omega} \langle w, v^t \rangle - g(w)$.
end for

step: if $g(\cdot)$ is not strongly convex, for example Lasso where $f_i(w) = \frac{1}{2}(y_i - w^T x_i)^2$ and $g(w) = \|w\|_1$, we redefine the formulation by adding a strongly convex term to $g(w)$ and subtract it from $f(w)$. More precisely, for $i = 1, \dots, n$, we define $\phi_i(w) = \frac{n+1}{2n} \|y_i - w^T x_i\|_2^2$, $\phi_{n+1}(w) = -\frac{\tilde{\lambda}(n+1)}{2} \|w\|_2^2$, $\tilde{g}(w) = \frac{1}{2} \|w\|_2^2 + \frac{\tilde{\lambda}}{\lambda} \|w\|_1$, and apply Algorithm 1 on $\frac{1}{n+1} \sum_{i=1}^{n+1} \phi_i(w) + \tilde{\lambda} \tilde{g}(w)$ (which is equivalent to Lasso), where the value of $\tilde{\lambda}$ will be specified later. Our analysis is thus focused on this alternative representation. The main challenge arises from the non-convexity of the newly defined ϕ_{n+1} , which precludes the use of the *dual method* as in Shalev-Shwartz and Zhang [2014]. While in Shalev-Shwartz [2016], a dual free SDCA algorithm is proposed and analyzed, the results do not apply to out setting for reasons discussed above.

Our contributions are two-fold. 1. We prove linear convergence of SDCA for a class of problem that are not strongly convex or even non-convex, making use of the concept *restricted strong convexity* (RSC). To the best of our knowledge, this is the first work to prove linear convergence of SDCA under this setting which includes several statistical model such as Lasso, Group Lasso, logistic regression, linear regression with SCAD regularization, corrected Lasso, to name a few. 2. As a byproduct, we derive a dual free from SDCA that extends the work of Shalev-Shwartz [2016] to account for more general regularization $g(w)$.

Related work. Agarwal et al. [2010] prove linear convergence of the *batched* gradient and composite gradient method in Lasso, Low rank matrix completion problem using the RSC condition. Loh and Wainwright [2013] extend this to a class of non-convex M-estimators. In spirit, our work can be thought as a stochastic counterpart. Recently, Qu et al. [2016, 2017] consider SVRG and SAGA in the similar setting as ours, but they look at the primal problem. Shalev-Shwartz [2016] considers dual free SDCA, but the analysis does not apply to the case where $F(\cdot)$ is not strongly convex. Similarly, Allen-Zhu and Yuan [2015b] consider the non-strongly convex setting in SVRG (SVRG⁺⁺), where $F(w)$ is convex, each individual $f_i(w)$ may be non-convex, but only establish sub-linear convergence. Recently, Li et al. [2016] consider SVRG with a zero-norm constraint and prove linear convergence for this specific formulation. In comparison, our results hold more generally, covering not only sparsity model but also corrected Lasso with noisy covariate, group sparsity model, and beyond.

2 Problem setup

In this paper we consider two setups, namely, (1) convex but not strongly convex $F(w)$ and (2) non-convex $F(w)$. For the convex case, we consider

$$\begin{aligned} \min_{g(w) \leq \rho} F(w) &:= f(w) + \lambda g(w) \\ &:= \frac{1}{n} \sum_{i=1}^n f_i(w) + \lambda g(w), \end{aligned} \quad (2)$$

where $\rho > 0$ is some pre-defined radius, $f_i(w)$ is the loss function for sample i , and $g(w)$ is a norm. Here we assume each $f_i(w)$ is convex and L_i -smooth.

For the non-convex $F(w)$ we consider

$$\begin{aligned} \min_{d_{\lambda}(w) \leq \rho} F(w) &:= f(w) + d_{\lambda, \mu}(w) \\ &= \frac{1}{n} \sum_{i=1}^n f_i(w) + d_{\lambda, \mu}(w), \end{aligned} \quad (3)$$

where $\rho > 0$ is some pre-defined radius, $f_i(w)$ is convex and L_i smooth, $d_{\lambda, \mu}(w)$ is a non-convex regularizer depending on a tuning parameter λ and a parameter μ explained in section 2.3. This M-estimator includes a side constraint depending on $d_{\lambda}(w)$, which needs to be a convex function and admits a lower bound $d_{\lambda}(w) \geq \|w\|_1$. It is also closed w.r.t. to $d_{\lambda, \mu}(w)$, more details are deferred to section 2.3.

We list some examples that belong to these two setups.

- Lasso: $F(w) = \sum_{i=1}^n \frac{1}{2}(y_i - w^T x_i)^2 + \lambda \|w\|_1$.
- Logistic regression with ℓ_1 penalty:

$$F(w) = \sum_{i=1}^n \log(1 + \exp(-y_i x_i^T w)) + \lambda \|w\|_1.$$

- Group Lasso $F(w) = \sum_{i=1}^n \frac{1}{2}(y_i - w^T x_i)^2 + \lambda \|w\|_{1,2}$.
- Corrected Lasso Loh and Wainwright [2011]: $F(w) = \sum_{i=1}^n \frac{1}{2n} (\langle w, x_i \rangle - y_i)^2 - \frac{1}{2} w^T \Sigma w + \lambda \|w\|_1$, where Σ is some positive definite matrix.
- linear regression with SCAD regularizer Fan and Li [2001]: $F(w) = \sum_{i=1}^n \frac{1}{2n} (\langle w, x_i \rangle - y_i)^2 + \text{SCAD}(w)$.

The first three examples belong to first setup while the last two belong to the second setup.

2.1 Restricted Strongly Convexity

Restricted Strongly Convexity (RSC) is proposed in Negahban et al. [2009], Van De Geer et al. [2009] and has been explored in several other work Agarwal et al. [2010], Loh and Wainwright [2013]. We say the loss function $f(w)$ satisfies the RSC condition with curvature κ and the tolerance parameter τ with respect to the norm $g(w)$ if

$$\begin{aligned} \Delta f(w_1, w_2) &\triangleq f(w_1) - f(w_2) - \langle \nabla f(w_2), w_1 - w_2 \rangle \\ &\geq \frac{\kappa}{2} \|w_1 - w_2\|_2^2 - \tau g^2(w_1 - w_2). \end{aligned} \quad (4)$$

When $f(w)$ is γ strongly convex, then it is RSC with $\kappa = \gamma$ and $\tau = 0$. However in many case, $f(w)$ may not be strongly convex, especially in the high dimensional case where the ambient dimension $p > n$. On the other hand, RSC is easier to be satisfied. Take Lasso as an example, under some mild condition, it is shown that Negahban et al. [2009] $\Delta f(w_1, w_2) \geq c_1 \|w_1 - w_2\|_2^2 - c_2 \frac{\log p}{n} \|w_1 - w_2\|_1^2$, where c_1, c_2 are some positive constants. Besides Lasso, the RSC condition holds for a large range of statistical models including log-linear model, group sparsity model, and low-rank model. See Negahban et al. [2009] for more detailed discussions.

2.2 Assumption on convex regularizer $g(w)$

Decomposability is the other ingredient needed to analyze the algorithm.

Definition: A regularizer is *decomposable* with respect to a pair of subspaces $\mathcal{A} \subseteq \mathcal{B}$ if

$$g(\alpha + \beta) = g(\alpha) + g(\beta) \quad \text{for all } \alpha \in \mathcal{A}, \beta \in \mathcal{B}^\perp,$$

where \perp means the orthogonal complement.

A concrete example is l_1 regularization for sparse vector supported on subset S . We define the subspace pairs with respect to the subset $S \subset \{1, \dots, p\}$, $\mathcal{A}(S) = \{w \in \mathbb{R}^p | w_j = 0 \text{ for all } j \notin S\}$ and $\mathcal{B}(S) = \mathcal{A}(S)^\perp$. The decomposability is thus easy to verify. Other widely used examples include non-overlap group norms such as $\|\cdot\|_{1,2}$, and the nuclear norm $\|\cdot\|_*$ Negahban et al. [2009]. In the rest of the paper, we denote $w_{\mathcal{B}}$ as the projection of w on the subspace \mathcal{B} .

Subspace compatibility constant

For any subspace \mathcal{A} of \mathbb{R}^p , the *subspace compatibility constant* with respect to the pair $(g, \|\cdot\|_2)$ is given by

$$\Psi(\mathcal{A}) = \sup_{u \in \mathcal{A} \setminus \{0\}} \frac{g(u)}{\|u\|_2}.$$

That is, it is the Lipschitz constant of the regularizer restricted in \mathcal{A} . For example, for the above-mentioned sparse vector with cardinality s , $\Psi(\mathcal{A}) = \sqrt{s}$ for $g(u) = \|u\|_1$.

2.3 Assumption on non-convex regularizer $d_{\lambda,\mu}(w)$

We consider regularizers that are *separable* across coordinates, i.e., $d_{\lambda,\mu}(w) = \sum_{j=1}^p \bar{d}_{\lambda,\mu}(w_j)$. Besides the separability, we make further assumptions on the univariate function $\bar{d}_{\lambda,\mu}(t)$:

1. $\bar{d}_{\lambda,\mu}(\cdot)$ satisfies $\bar{d}_{\lambda,\mu}(0) = 0$ and is symmetric about zero (i.e., $\bar{d}_{\lambda,\mu}(t) = \bar{d}_{\lambda,\mu}(-t)$).
2. On the nonnegative real line, $\bar{d}_{\lambda,\mu}(\cdot)$ is nondecreasing.
3. For $t > 0$, $\frac{\bar{d}_{\lambda,\mu}(t)}{t}$ is nonincreasing in t .
4. $\bar{d}_{\lambda,\mu}(\cdot)$ is differentiable for all $t \neq 0$ and subdifferentiable at $t = 0$, with $\lim_{t \rightarrow 0^+} \bar{d}'_{\lambda,\mu}(t) = \lambda L_d$.
5. $\bar{d}_\lambda(t) := (\bar{d}_{\lambda,\mu}(t) + \frac{\mu}{2}t^2)/\lambda$ is convex.

For instance, SCAD satisfying these assumptions.

$$\mathbf{SCAD}_{\lambda, \zeta}(\mathbf{t}) =$$

$$\begin{cases} \lambda|t|, & \text{for } |t| \leq \lambda, \\ -(t^2 - 2\zeta\lambda|t| + \lambda^2)/(2(\zeta - 1)), & \text{for } \lambda < |t| \leq \zeta\lambda, \\ (\zeta + 1)\lambda^2/2, & \text{for } |t| > \zeta\lambda, \end{cases}$$

where $\zeta > 2$ is a fixed parameter. It satisfies the assumption with $L_d = 1$ and $\mu = \frac{1}{\zeta-1}$ Loh and Wainwright [2013].

2.4 Applying SDCA

2.4.1 Convex $F(w)$

Following a similar line as we did for Lasso, to apply the SDCA algorithm, we define $\phi_i(w) = \frac{n+1}{n}f_i(w)$ for $i = 1, \dots, n$, $\phi_{n+1}(w) = -\frac{\tilde{\lambda}(n+1)}{2}\|w\|_2^2$, $\tilde{g}(w) = \frac{1}{2}\|w\|_2^2 + \frac{\lambda}{\tilde{\lambda}}g(w)$. Correspondingly, the new smoothness parameters are $\tilde{L}_i = \frac{n+1}{n}L_i$ for $i = 1, \dots, n$ and $\tilde{L}_{n+1} = \tilde{\lambda}(n+1)$. Problem (2) is thus equivalent to the following

$$\min_{w \in \Omega} \frac{1}{1+n} \sum_{i=1}^{n+1} \phi_i(w) + \tilde{\lambda} \tilde{g}(w). \quad (5)$$

This enables us to apply Algorithm 1 to the problem with $\phi_i, \tilde{g}, \tilde{L}_i, \tilde{\lambda}$. In particular, while ϕ_{n+1} is not convex, \tilde{g} is still convex (1-strongly convex). We exploit this property in the proof and define $\tilde{g}^*(v) = \max_{w \in \Omega} \langle w, v \rangle - \tilde{g}(w)$, where Ω is a convex compact set. Since $\tilde{g}(w)$ is 1-strongly convex, $\tilde{g}^*(v)$ is 1-smooth [Theorem 1 of Nesterov, 2005].

2.4.2 Non-convex $F(w)$

Similarly, we define

$$\begin{aligned} \phi_i(w) &= \frac{n+1}{n}f_i(w) \quad \text{for } i = 1, \dots, n, \\ \phi_{n+1}(w) &= -\frac{(\tilde{\lambda} + \mu)(n+1)}{2}\|w\|_2^2, \\ \tilde{g}(w) &= \frac{1}{2}\|w\|_2^2 + \frac{\lambda}{\tilde{\lambda}}d_\lambda(w), \end{aligned}$$

and then apply Algorithm 1 on

$$\min_{w \in \Omega} \frac{1}{1+n} \sum_{i=1}^{n+1} \phi_i(w) + \tilde{\lambda} \tilde{g}(w).$$

The update rule of proximal step for different $g(\cdot), d_\lambda(\cdot)$ (such as SCAD and MCP) can be found in Loh and Wainwright [2013].

3 Theoretical Result

In this section, we present the main theoretical results, and some corollaries that instantiate the main results in several well known statistical models.

3.1 Convex $F(w)$

To begin with, we define several terms related to the algorithm.

- w^* is true unknown parameter. $g^*(\cdot)$ is the dual norm of $g(\cdot)$. Conjugate function $\tilde{g}^*(v) = \max_{w \in \Omega} \langle w, v \rangle - \tilde{g}(w)$, where $\Omega = \{w | g(w) \leq \rho\}$.
- \hat{w} is the optimal solution to problem 2, and we assume \hat{w} is in the interior of Ω w.l.o.g. by choosing Ω large enough.
- (\hat{w}, \hat{v}) is an optimal solution pair satisfying $\tilde{g}(\hat{w}) + \tilde{g}^*(\hat{v}) = \langle \hat{w}, \hat{v} \rangle$.
- $A_t = \sum_{j=1}^{n+1} \frac{1}{q_j} \|a_j^t - \hat{a}_j\|_2^2$, where $\hat{a}_j = -\nabla \phi_j(\hat{w})$. $B_t = 2(\tilde{g}^*(v^t) - \langle \nabla \tilde{g}^*(\hat{v}), v^t - \hat{v} \rangle - \tilde{g}^*(\hat{v}))$.

We remark that our definition on the first potential A_t is the same as in Shalev-Shwartz [2016], while the second one B_t is different. If $\tilde{g}(w) = \frac{1}{2}\|w\|_2^2$, our definition on B_t reduces to that in Shalev-Shwartz [2016], i.e., $\|w^t - \hat{w}\|_2^2$. To see this, when $\tilde{g}(w) = \frac{1}{2}\|w\|_2^2$, $\tilde{g}^*(v) = \frac{1}{2}\|v\|_2^2$ and $v^t = w^t, \hat{v} = \hat{w}$. We then define another potential C_t , which is a combination of A_t and B_t . $C_t = \frac{\eta}{(n+1)^2} A_t + \frac{\tilde{\lambda}}{2} B_t$. Notice that using smoothness of $\tilde{g}^*(\cdot)$ and Lemma 1 and 2 in the appendix, it is not hard to show $B_t \geq \|w^t - \hat{w}\|_2^2$. Thus if C_t converges, so does $\|w^t - \hat{w}\|_2^2$.

For notational simplicity, we define the following two terms used in the theorem.

- Effective RSC parameter: $\tilde{\kappa} = \kappa - 64\tau\Psi^2(\mathcal{B})$.
- Tolerance: $\delta = c_1\tau\delta_{stat}^2$, where $\delta_{stat}^2 = (\Psi(\mathcal{B})\|\hat{w} - w^*\|_2 + g(w_{\mathcal{A}^\perp}^*))^2$, c_1 is a universal positive constant.

Theorem 1. Assume each $f_i(w)$ is L_i smooth and convex, $f(w)$ satisfies the RSC condition with parameter (κ, τ) , w^* is feasible, the regularizer is decomposable w.r.t $(\mathcal{A}, \mathcal{B})$ such that $\tilde{\kappa} \geq 0$, and the Algorithm 1 runs with $\eta \leq \min\{\frac{1}{16(\tilde{\lambda} + \bar{L})}, \frac{1}{4\tilde{\lambda}(n+1)}\}$, where $\bar{L} = \frac{1}{n} \sum_{i=1}^n L_i$, $\tilde{\lambda}$ is chosen such that $0 < \tilde{\lambda} \leq \tilde{\kappa}$. If we choose the regularization parameter λ such that $\lambda \geq \max(2g^*(\nabla f(w^*)), c\tau\rho)$ where c is some universal positive constant, then we have

$$\mathbb{E}(C_t) \leq (1 - \eta\tilde{\lambda})^t C_0,$$

until $F(w^t) - F(\hat{w}) \leq \delta$, where the expectation is for the randomness of sampling of i in the algorithm.

Some remarks are in order for interpreting the theorem.

1. In several statistical models, the requirement of $\tilde{\kappa} > 0$ is easy to satisfy under mild conditions. For instance, in Lasso we have $\Psi^2(\mathcal{B}) = s$. $\tau \simeq \frac{\log p}{n}$, $\kappa = \frac{1}{2}$ if the feature vector x_i is sampled from $N(0, I)$. Thus, if $\frac{128s \log p}{n} \leq 1$, we have $\tilde{\kappa} \geq \frac{1}{4}$.
2. In some models, we can choose the pair of subspace $(\mathcal{A}, \mathcal{B})$ such that $w^* \in \mathcal{A}$ and thus $\delta = c_1\tau\Psi^2(\mathcal{B})\|\hat{w} - w^*\|_2^2$. In Lasso, as we mentioned above $\tau \simeq \frac{\log p}{n}$, thus $\delta \simeq c_1s\frac{\log p}{n}\|\hat{w} - w^*\|_2^2$, i.e., this tolerance is dominated by statistical error if s is small and n is large.
3. We know $B_t \geq \|w^t - \hat{w}\|_2^2$ and $A_t \geq 0$, thus $C_t \geq \frac{\tilde{\lambda}}{2}\|w^t - \hat{w}\|_2^2$, thus, if C_t converge, so does $\|w^t - \hat{w}\|_2^2$. When $F(w^t) - F(\hat{w}) \leq \delta$, using Lemma 5 in the supplementary material, it is easy to get $\|w^t - \hat{w}\|_2^2 \leq \frac{4\delta}{\tilde{\kappa}}$.

Combining these remarks, Theorem 1 states that the optimization error decreases geometrically until it achieves the tolerance $\delta/\tilde{\kappa}$ which is dominated by the statistical error $\|\hat{w} - w^*\|_2^2$, thus can be ignored from the statistical view.

If $g(w)$ in Problem (1) is indeed 1-strongly convex, we have the following proposition, which extends dual-free SDCA Shalev-Shwartz [2016] into the general regularization case. Notice we now directly apply the Algorithm 1 on Problem (1) and change the definitions of A_t and B_t correspondingly. In particular, $A_t = \sum_{j=1}^n \frac{1}{q_j} \|a_j^t - \hat{a}_j\|_2^2$, $B_t = 2(g^*(v^t) - \langle \nabla g^*(\hat{v}), v^t - \hat{v} \rangle - g^*(\hat{v}))$, where $\hat{a}_j = -\nabla f_j(\hat{w})$. C_t is still defined in the same way, i.e., $C_t = \frac{\eta}{n^2} A_t + \frac{\lambda}{2} B_t$.

Proposition 1. *Suppose each $f_i(w)$ is L_i smooth, $f(w)$ is convex, $g(w)$ is 1- strongly convex, \hat{w} is the optimal solution of Problem (1), and $\bar{L} = \frac{1}{n} \sum_{i=1}^n L_i$. Then we have: (I) If each $f_i(x)$ is convex, we run the Algorithm 1 with $\eta \leq \min\{\frac{1}{4L}, \frac{1}{4\lambda n}\}$, then $\mathbb{E}(C_t) \leq (1 - \eta\lambda)^t C_0$. (II) Otherwise, we run the Algorithm 1 with $\eta \leq \min\{\frac{\lambda}{4L^2}, \frac{1}{4\lambda n}\}$, and $\mathbb{E}(C_t) \leq (1 - \eta\lambda)^t C_0$. Note that $B_t \geq \|w^t - \hat{w}\|_2^2$, $A_t \geq 0$, thus $\|w^t - \hat{w}\|_2^2$ decreases geometrically.*

In the following we present several corollaries that instantiate Theorem 1 with several concrete statistical models. This essentially requires to choose appropriate subspace pair $(\mathcal{A}, \mathcal{B})$ in these models and check the RSC condition.

3.1.1 Sparse linear regression

Our first example is Lasso, where $f_i(w) = \frac{1}{2}(y_i - w^T x_i)^2$ and $g(w) = \|w\|_1$. We assume each feature vector x_i is generated from Normal distribution $N(0, \Sigma)$ and the true parameter $w^* \in \mathbb{R}^p$ is sparse with cardinality s . The observation y_i is generated by $y_i = (w^*)^T x_i + \xi_i$, where ξ_i is a Gaussian noise with mean 0 and variance σ^2 . We denote the data matrix by $X \in \mathbb{R}^{n \times p}$ and X_j is the j th column of X . Without loss of generality, we assume X is column normalized, i.e., $\frac{\|X_j\|_2}{\sqrt{n}} \leq 1$ for all $j = 1, 2, \dots, p$. We denote $\sigma_{\min}(\Sigma)$ as the smallest eigenvalue of Σ , and $\nu(\Sigma) = \max_{i=1,2,\dots,p} \Sigma_{ii}$.

Corollary 1. *Assume w^* is the true parameter supported on a subset with cardinality at most s , and we choose the parameter $\lambda, \tilde{\lambda}$ such that $\lambda \geq \max(6\sigma\sqrt{\frac{\log p}{n}}, c_1\rho\nu(\Sigma)\frac{\log p}{n})$ and $0 < \tilde{\lambda} \leq \tilde{\kappa}$ hold, where $\tilde{\kappa} = \frac{1}{2}\sigma_{\min}(\Sigma) - c_2\nu(\Sigma)\frac{s\log p}{n}$, where c_1, c_2 are some universal positive constants. Then we run the Algorithm 1 with $\eta \leq \min\{\frac{1}{16(\lambda+\tilde{\lambda})}, \frac{1}{4\lambda(n+1)}\}$ and have*

$$\mathbb{E}(C_t) \leq (1 - \eta\tilde{\lambda})^t C_0$$

with probability at least $1 - \exp(-3\log p) - \exp(-c_3 n)$, until $F(w^t) - F(\hat{w}) \leq \delta$, where $\delta = c_4\nu(\Sigma)\frac{s\log p}{n}\|\hat{w} - w^\|_2^2$. c_1, c_2, c_3, c_4 are some universal positive constants.*

The requirement $\lambda \geq 6\sigma\sqrt{\frac{\log p}{n}}$ is documented in literature Negahban et al. [2009] to ensure that Lasso is statistically consistent. And $\lambda \geq c_1 \rho\nu(\Sigma)\frac{\log p}{n}$ is needed for fast convergence of optimization algorithms, which is similar to the condition proposed in in Agarwal et al. [2010] for batch optimization algorithm. When $\frac{s\log p}{n} = o(1)$, which is necessary for statistical consistency of Lasso, we have $\frac{1}{2}\sigma_{\min}(\Sigma) - c_2\nu(\Sigma)\frac{s\log p}{n} \geq 0$, which guarantees the existence of $\tilde{\lambda}$. Also notice under this condition, $\delta = c_3\nu(\Sigma)\frac{s\log p}{n}\|\hat{w} - w^*\|_2^2$ is of a lower order of $\|\hat{w} - w^*\|_2^2$. Using remark 3 in Theorem 1, we have $\|w^t - \hat{w}\|_2^2 \leq \frac{4\delta}{\tilde{\kappa}}$, which is dominated by the statistical error $\|\hat{w} - w^*\|_2^2$ and hence can be ignored from the statistical perspective. Thus to sum up, Corollary 1 states the optimization error decreases geometrically until it achieves the statistical limit of Lasso.

3.1.2 Group Sparsity Model

Yuan and Lin [2006] introduce the group Lasso to allow predefined groups of covariates to be selected together into or out of a model together. The most commonly used regularizer to encourage group sparsity is $\|\cdot\|_{1,2}$. In the following, we define group sparsity formally. We assume groups are disjointed, i.e., $\mathcal{G} = \{G_1, G_2, \dots, G_{N_G}\}$ and $G_i \cap G_j = \emptyset$. The regularization is $\|w\|_{\mathcal{G},q} \triangleq \sum_{g=1}^{N_G} \|w_g\|_q$. When $q = 2$, it reduces to the commonly used group Lasso Yuan and Lin [2006], and another popularly used case is $q = \infty$ Turlach et al. [2005], Quattoni et al. [2009]. We require the following condition, which generalizes the column normalization condition in the Lasso case. Given a group G of size m and $X_G \in \mathbb{R}^{n \times m}$, the associated operator norm $\|X_G\|_{q \rightarrow 2} \triangleq \max_{\|w\|_q=1} \|X_G w\|_2$ satisfies

$$\frac{\|X_G\|_{q \rightarrow 2}}{\sqrt{n}} \leq 1 \quad \text{for all } i = 1, 2, \dots, N_G.$$

The condition reduces to the column normalized condition when each group contains only one feature (i.e., Lasso).

We now define the subspace pair $(\mathcal{A}, \mathcal{B})$ in the group sparsity model. For a subset $S_G \subseteq \{1, \dots, N_G\}$ with cardinality $s_G = |S_G|$, we define the subspace

$$\mathcal{A}(S_G) = \{w | w_{G_i} = 0 \text{ for all } i \notin S_G\},$$

and $\mathcal{A} = \mathcal{B}$. The orthogonal complement is

$$\mathcal{B}^\perp(S_G) = \{w | w_{G_i} = 0 \text{ for all } i \in S_G\}.$$

We can easily verify that

$$\|\alpha + \beta\|_{\mathcal{G},q} = \|\alpha\|_{\mathcal{G},q} + \|\beta\|_{\mathcal{G},q},$$

for any $\alpha \in \mathcal{A}(S_G)$ and $\beta \in \mathcal{B}^\perp(S_G)$.

In the following corollary, we use $q = 2$, i.e., group Lasso, as an example. We assume the observation y_i is generated by $y_i = x_i^T w^* + \xi_i$, where $x_i \sim N(0, \Sigma)$, and $\xi_i \sim N(0, \sigma^2)$.

Corollary 2 (Group Lasso). *Assume $w \in \mathbb{R}^p$ and each group has m parameters, i.e., $p = mN_G$. Denote by s_G the cardinality of non-zero group, and we choose parameters $\lambda, \tilde{\lambda}$ such that*

$$\begin{aligned} \lambda &\geq \max \left(4\sigma \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{\log N_G}{n}} \right), c_1 \rho \sigma_2(\Sigma) \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{3 \log N_G}{n}} \right)^2 \right); \\ \text{and } 0 &< \tilde{\lambda} \leq \tilde{\kappa}, \quad \text{where } \tilde{\kappa} = \sigma_1(\Sigma) - c_2 \sigma_2(\Sigma) s_G \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{3 \log N_G}{n}} \right)^2; \end{aligned}$$

where $\sigma_1(\Sigma)$ and $\sigma_2(\Sigma)$ are positive constant depending only on Σ . If we run the Algorithm 1 with $\eta \leq \min\{\frac{1}{16(\tilde{\lambda}+L)}, \frac{1}{4\tilde{\lambda}(n+1)}\}$, then we have

$$\mathbb{E}(C_t) \leq (1 - \eta \tilde{\lambda})^t C_0$$

with probability at least $1 - 2 \exp(-2 \log N_G) - c_2 \exp(-c_3 n)$, until $F(w^t) - F(\hat{w}) \leq \delta$, where $\delta = c_3 \sigma_2(\Sigma) s_G \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{3 \log N_G}{n}} \right)^2 \|\hat{w} - w^*\|_2^2$.

We offer some discussions to interpret the corollary. To satisfy the requirement $\tilde{\kappa} \geq 0$, it suffices to have

$$s_{\mathcal{G}} \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{3 \log N_{\mathcal{G}}}{n}} \right)^2 = o(1).$$

This is a mild condition, as it is needed to guarantee the statistical consistency of group Lasso Negahban et al. [2009]. Notice that the condition is easily satisfied when $s_{\mathcal{G}}$ and m are small. Under this same condition, since $\delta = c_3 \sigma_2(\Sigma) s_{\mathcal{G}} \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{3 \log N_{\mathcal{G}}}{n}} \right)^2 \|\hat{w} - w^*\|_2^2$, we conclude that δ is dominated by $\|\hat{w} - w^*\|_2^2$. Again, it implies the optimization error decrease geometrically up to the scale $o(\|\hat{w} - w^*\|_2^2)$ which is dominated by the statistical error of the model.

3.1.3 Extension to generalized linear model

We consider the *generalized linear model* of the following form,

$$\min_{w \in \Omega} \frac{1}{n} \sum_{i=1}^n (\Phi(w, x_i) - y_i \langle w, x_i \rangle) + \lambda \|w\|_1,$$

which covers such case as Lasso (where $\Phi(\theta) = \frac{\theta^2}{2}$) and logistic regression (where $\Phi(\theta) = \log(1 + \exp(\theta))$). In this model, we have

$$\Delta f(w_1, w_2) = \frac{1}{n} \sum_{i=1}^n \Phi''(\langle w_t, x_i \rangle) \langle x_i, w_1 - w_2 \rangle^2,$$

where $w_t = tw_1 + (1-t)w_2$ for some $t \in [0, 1]$. The RSC condition thus is equivalent to:

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \Phi''(\langle w_t, x_i \rangle) \langle x_i, w_1 - w_2 \rangle^2 \\ & \geq \frac{\kappa}{2} \|w_1 - w_2\|_2^2 - \tau g^2(w_1 - w_2) \quad \text{for } w_1, w_2 \in \Omega. \end{aligned} \tag{6}$$

Here we require Ω to be a bounded set Loh and Wainwright [2013]. This requirement is essential since in some generalized linear model $\Phi''(\theta)$ approaches to zero as θ diverges. For instance, in logistic regression, $\Phi''(\theta) = \exp(\theta)/(1 + \exp(\theta))^2$, which tends to zero as $\theta \rightarrow \infty$. For a broad class of generalized linear models, RSC holds with $\tau = c \frac{\log p}{n}$, thus the same result as that of Lasso holds, modulus change of constants.

3.2 Non-convex $F(w)$

In the non-convex case, we assume the following RSC condition:

$$\Delta f(w_1, w_2) \geq \frac{\kappa}{2} \|w_1 - w_2\|_2^2 - \tau \|w_1 - w_2\|_1^2$$

with $\tau = \theta \frac{\log p}{n}$ for some constant θ . We again define the potential A_t, B_t and C_t in the same way with convex case. The main difference is that now we have $\phi_{n+1} = -\frac{(\tilde{\lambda} + \mu)(n+1)}{2} \|w\|_2^2$ and the effective RSC parameter $\tilde{\kappa}$ is different. The necessary notations for presenting the theorem are listed below:

- w^* is the unknown true parameter that is s -sparse. Conjugate function $\tilde{g}^*(v) = \max_{w \in \Omega} \langle w, v \rangle - \tilde{g}(w)$, where $\Omega = \{w | d_\lambda(w) \leq \rho\}$. Note Ω is convex due to convexity of $d_\lambda(w)$.
- \hat{w} is the global optimum of Problem (3), we assume it is in the interior of Ω w.l.o.g.
- (\hat{w}, \hat{v}) is an optimal solution pair satisfying $\tilde{g}(\hat{w}) + \tilde{g}^*(\hat{v}) = \langle \hat{w}, \hat{v} \rangle$.
- $A_t = \sum_{j=1}^{n+1} \frac{1}{q_j} \|a_j^t - \hat{a}_j\|_2^2$, $B_t = 2(\tilde{g}^*(v^t) - \langle \nabla \tilde{g}^*(\hat{v}), v^t - \hat{v} \rangle - \tilde{g}^*(\hat{v}))$, $C_t = \frac{\eta}{(n+1)^2} A_t + \frac{\tilde{\lambda}}{2} B_t$, where $\hat{a}_j = -\nabla \phi_j(\hat{w})$.
- Effective RSC parameter: $\tilde{\kappa} = \kappa - \mu - 64\tau s$, where $\tau = \theta \frac{\log p}{n}$ for some constant θ . Tolerance: $\delta = c_1 \tau s \|\hat{w} - w^*\|_2^2$, where c_1 is a universal positive constant.

Theorem 2. Suppose w^* is s sparse, \hat{w} is the global optimum of Problem (3). Assume each $f_i(w)$ is L_i smooth and convex, $f(w)$ satisfies the RSC condition with parameter (κ, τ) , where $\tau = \theta \frac{\log p}{n}$ for some constant θ , $d_{\lambda, \mu}(w)$ satisfies the Assumption in section 2.3, $\lambda L_d \geq \max\{c\rho\theta \frac{\log p}{n}, 4\|\nabla f(w^*)\|_\infty\}$ where c is some universal positive constant, and $\tilde{\kappa} - \mu \geq \tilde{\lambda} > 0$, the Algorithm 1 runs with $\eta \leq \min\{\frac{1}{16(\tilde{\lambda} + \bar{L})}, \frac{1}{4\lambda(n+1)}\}$, where $\bar{L} = \frac{1}{n} \sum_{i=1}^n L_i$, then we have

$$\mathbb{E}(C_t) \leq (1 - \eta \tilde{\lambda})^t C_0,$$

until $F(w^t) - F(\hat{w}) \leq \delta$, where the expectation is for the randomness of sampling of i in the algorithm.

Some remarks are in order to interpret the theorem.

- We require $\tilde{\lambda}$ to satisfy $0 < \tilde{\lambda} \leq \kappa - 2\mu - 64\theta \frac{\log p}{n} s$. Thus if the non-convex parameter μ is too large, we can not find such $\tilde{\lambda}$.
- Note that $\delta = c_1 s \theta \frac{\log p}{n} \|\hat{w} - w^*\|_2^2$ is dominated by $\|\hat{w} - w^*\|_2^2$ when the model is sparse and n is large. Similar as the convex case, by $B_t \geq \|w^t - \hat{w}\|_2^2$, the theorem says the optimization error decreases linearly up to the fundamental statistical error of the model.

The first non-convex model we consider is linear regression with SCAD. Here $f_i(w) = \frac{1}{2}(\langle w, x_i \rangle - y_i)^2$ and $d_{\lambda, \mu}(\cdot)$ is $SCAD(\cdot)$ with parameter λ and ζ . The data (x_i, y_i) are generated similarly as in the Lasso case.

Corollary 3 (Linear regression with SCAD regularization). Suppose we have n i.i.d. observations $\{(x_i, y_i)\}$, w^* is s sparse, \hat{w} is global optima and we choose λ and $\tilde{\lambda}$ such that $\lambda \geq \max\{c_1 \rho \nu(\Sigma) \frac{\log p}{n}, 12\sigma \sqrt{\frac{\log p}{n}}\}$, $\tilde{\kappa} - \frac{1}{\zeta-1} \geq \tilde{\lambda} > 0$, then we run the algorithm with $\eta \leq \min\{\frac{1}{16(\tilde{\lambda} + \bar{L})}, \frac{1}{4\lambda(n+1)}\}$, where $\bar{L} = \frac{1}{n} \sum_{i=1}^n L_i$, then have

$$\mathbb{E}(C_t) \leq (1 - \eta \tilde{\lambda})^t C_0,$$

with $1 - \exp(-3 \log p) - \exp(-c_2 n)$, until $F(w^t) - F(\hat{w}) \leq \delta$, where $\tilde{\kappa} = \frac{1}{2} \sigma_{\min}(\Sigma) - c_3 \nu(\Sigma) \frac{s \log p}{n} - \frac{1}{\zeta-1}$, $\delta = c_4 \nu(\Sigma) \frac{s \log p}{n} \|\hat{w} - w^*\|_2^2$. Here c_1, c_2, c_3, c_4 are universal positive constants.

If the non-convex parameter $\frac{1}{\zeta-1}$ is small, s is sparse, n is large, then we can choose a positive $\tilde{\lambda}$ to guarantee the convergence of algorithm. Under this setting, we have the tolerance δ dominated by $\|\hat{w} - w^*\|_2^2$.

The second example is the *corrected Lasso*. In many applications, the covariate may be observed subject to corruptions. In this section, we consider corrected Lasso proposed by Loh and Wainwright [2011]. Suppose data are generated according to a linear model $y_i = x_i^T w^* + \xi_i$, where ξ_i is a random zero-mean sub-Gaussian noise with variance σ^2 ; and each data point x_i is i.i.d. sampled from a zero-mean normal distribution, i.e., $x_i \sim N(0, \Sigma)$. We denote the data matrix by $X \in \mathbb{R}^{n \times p}$, the smallest eigenvalue of Σ by $\sigma_{\min}(\Sigma)$, and the largest eigenvalue of Σ by $\sigma_{\max}(\Sigma)$.

The observation z_i of x_i is corrupted by additive noise, in particular, $z_i = x_i + \varsigma_i$, where $\varsigma_i \in \mathbb{R}^p$ is a random vector independent of x_i , say zero-mean with known covariance matrix Σ_ς . Define $\hat{\Gamma} = \frac{Z^T Z}{n} - \Sigma_\varsigma$ and $\hat{\gamma} = \frac{Z^T y}{n}$. Our goal is to estimate w^* based on y_i and z_i (but not x_i which is not observable), and the corrected Lasso proposes to solve the following:

$$\hat{w} \in \arg \min_{\|w\|_1 \leq \rho} \frac{1}{2} w^T \hat{\Gamma} w - \hat{\gamma} w + \lambda \|w\|_1.$$

Equivalently, it solves

$$\min_{\|w\|_1 \leq \rho} \frac{1}{2n} \sum_{i=1}^n (y_i - w^T z_i)^2 - \frac{1}{2} w^T \Sigma_\varsigma w + \lambda \|w\|_1.$$

Notice that due to the term $-\frac{1}{2} w^T \Sigma_\varsigma w$, the optimization problem is non-convex.

Corollary 4 (Corrected Lasso). *Suppose we are given i.i.d. observation $\{(z_i, y_i)\}$ from the linear model with additive noise, w^* is s sparse, and $\Sigma_\varsigma = \gamma_\varsigma I$. Let \hat{w} be the global optima. We choose $\lambda \geq \max\{c_0 \rho \frac{\log p}{n}, c_1 \varphi \sqrt{\frac{\log p}{n}}\}$, and positive $\tilde{\lambda}$ such that $\tilde{\kappa} - \gamma_\varsigma \geq \tilde{\lambda} > 0$, where $\varphi = (\sqrt{\sigma_{\max}(\Sigma)} + \sqrt{\gamma_\varsigma})(\sigma + \sqrt{\gamma_\varsigma} \|w^*\|_2)$, and*

$$\tilde{\kappa} = \frac{1}{2} \sigma_{\min}(\Sigma) - c_2 \sigma_{\min}(\Sigma) \max \left(\left(\frac{\sigma_{\max}(\Sigma) + \gamma_\varsigma}{\sigma_{\min}(\Sigma)} \right)^2, 1 \right) \frac{s \log p}{n} - \gamma_\varsigma,$$

then if we run the algorithm with $\eta \leq \min\{\frac{1}{16(\tilde{\lambda} + L)}, \frac{1}{4\tilde{\lambda}(n+1)}\}$, we have

$$\mathbb{E}(C_t) \leq (1 - \eta \tilde{\lambda})^t C_0,$$

with probability at least $1 - c_3 \exp \left(-c_4 n \min \left(\frac{\sigma_{\min}^2(\Sigma)}{(\sigma_{\max}(\Sigma) + \gamma_\varsigma)^2}, 1 \right) \right) - c_5 \exp(-c_6 \log p)$, until $F(w^t) - F(\hat{w}) \leq \delta$, where $\delta = c_7 \sigma_{\min}(\Sigma) \max \left(\left(\frac{\sigma_{\max}(\Sigma) + \gamma_\varsigma}{\sigma_{\min}(\Sigma)} \right)^2, 1 \right) \frac{s \log p}{n} \|\hat{w} - w^*\|_2^2$, and c_0 to c_7 are some universal positive constants.

Some remarks of the corollary are in order.

- The result can be easily extended to more general $\Sigma_\varsigma \preceq \gamma_\varsigma I$.
- The requirement of λ is similar with its batch counterpart Loh and Wainwright [2011].
- Similar to the setting in Lasso, we need $\frac{s \log p}{n} = o(1)$. To ensure the existence of such $\tilde{\lambda}$, the non-convex parameter γ_ς can not be too large, which is similar to the result in Loh and Wainwright [2013].

4 Experiment Result

We report numerical experiment results to validate our theoretical findings, namely, without strong convexity, SDCA still achieves linear convergence under our setup. The setup of experiment is similar to that in Qu et al. [2016]. On both synthetic and real datasets we report results of SDCA and compare with several other algorithms. These algorithms are Prox-SVRG Xiao and Zhang [2014], SAGA Defazio et al. [2014], Prox-SAG an proximal version of the algorithm in Schmidt et al. [2013], proximal stochastic gradient (Prox-SGD), regularized dual averaging method (RDA) Xiao [2010] and the proximal full gradient method (Prox-GD) Nesterov [2013]. For the algorithms with a constant learning rate (i.e., Prox-SAG, SAGA, Prox-SVRG, SDCA, Prox-GD), we tune the learning rate from an exponential grid $\{2, \frac{2}{2^1}, \dots, \frac{2}{2^{12}}\}$ and chose the one with best performance. Below are some further remarks.

- Convergence rates of Prox-SGD and RDA are sub-linear according to the analysis in Bottou [2010], Xiao [2010]. The stepsize of Prox-SGD is set as $\eta_k = \eta_0/\sqrt{k}$ as suggested in Duchi and Singer [2009] and that of RDA is $\beta_k = \beta_0\sqrt{k}$ as suggested in Xiao [2010]. η_0 and β_0 are chosen as the one that attains the best performance among powers of 10.
- Agarwal et al. [2010] prove that Prox-GD converges linearly in the setting of our experiment.
- The convergence rate of Prox-SVRG and SAGA is linear in our setting, shown recently in Qu et al. [2016, 2017].
- To the best of our knowledge, linear convergence of Prox-SAG in our setting has not been established, although in practice it works well.


4.1 Synthetic dataset

In this section we test algorithms above on Lasso, Group Lasso, Corrected Lasso and SCAD .


4.1.1 Lasso

We generated the feature vector $x_i \in \mathbb{R}^p$ independently from $N(0, \Sigma)$, where we set $\Sigma_{ii} = 1$, for $i = 1, \dots, p$ and $\Sigma_{ij} = b$, for $i \neq j$. The responds y_i is generated by follows: $y_i = x_i^T w^* + \xi_i$. $w^* \in \mathbb{R}^p$ is a sparse vector with cardinality s , where the non-zero entries are ± 1 drawn from the Bernoulli distribution with probability 0.5. The noise ξ_i follows the standard normal distribution. We set $\lambda = 0.05$ in Lasso and choose $\tilde{\lambda} = 0.25$ in SDCA. In the following experiment, we set $p = 5000$, $n = 2500$ and try different settings on s and b .


Figure 1 presents simulation results on Lasso. Among all four settings, SDCA, SVRG and SAG converge with linear rates. When $b = 0.4$ and $s = 100$, SDCA outperforms the other two. When $b = 0$, the Prox-GD works although with slower rate. While b is nonzero, the Prox-GD does not works well due to the large condition number. RDA and SGD converges slowly in all setting because of the large variance in gradient.




(a) $s=50, b=0$



(b) $s=100, b=0$



(c) $s=50, b=0.1$



(d) $s=100, b=0.4$

Figure 1: Comparison between six algorithms on Lasso. The x-axis is the number of pass over the dataset. y-axis is the objective gap $F(w_k) - F(\hat{w})$ with a log scale. In figure (a), $s=50, b=0$. In figure (b), $s=100, b=0$. In figure (c), $s=50, b=0.1$. In figure (d), $s=100, b=0.4$.




Figure 2: Comparison between six algorithms on group Lasso. The x-axis is the number of pass over the dataset. y-axis is the objective gap $F(w_k) - F(\hat{w})$ with a log scale. In Figure (a), $m=10, s_G = 10, b = 0$. In Figure (b), $m=20, s_G = 20, b = 0$. In Figure (c), $m=10, s_G = 10, b = 0.1$. In Figure (d), $m=20, s_G = 20, b = 0.4$.

4.1.2 Group Lasso

We report the experiment result on Group Lasso in Figure 2. Similarly as Lasso, we generate the observation $y_i = x_i^T w^* + \xi_i$ with the feature vectors independently sampled from $N(0, \Sigma)$, where $\Sigma_{ii} = 1$ and $\Sigma_{ij} = b, i \neq j$. The cardinality of non-zero group is s_G , and the non-zero entries are sampled uniformly from $[-1, 1]$. In the following experiment, we try different setting on b , group size m and group sparsity s_G .

Similar with the result in Lasso, SDCA, SVRG and SAG performs well in all settings. When $m = 20$ and $s_G = 20$, SDCA outperforms the other two. The Prox-GD converges with linear rate when $m = 10, s_G = 10, b = 0$ but does not work well in the other three setting. SGD and RDA converge slowly in all four settings.

4.1.3 Corrected Lasso

We generate data as follows: $y_i = x_i^T w^* + \xi_i$, where each data point $x_i \in \mathbb{R}^p$ is drawn from normal distribution $N(0, I)$, and the noise ξ_i is drawn from $N(0, 1)$. The coefficient w^* is sparse with cardinality s , where the non-zero coefficient equals to ± 1 generated from

the Bernoulli distribution with probability 0.5. We set covariance matrix $\Sigma_\zeta = \gamma_\zeta I$. We choose $\lambda = 0.05$ in the formulation and $\tilde{\lambda} = 0.1$ in SDCA. The result is presented in Figure 3.

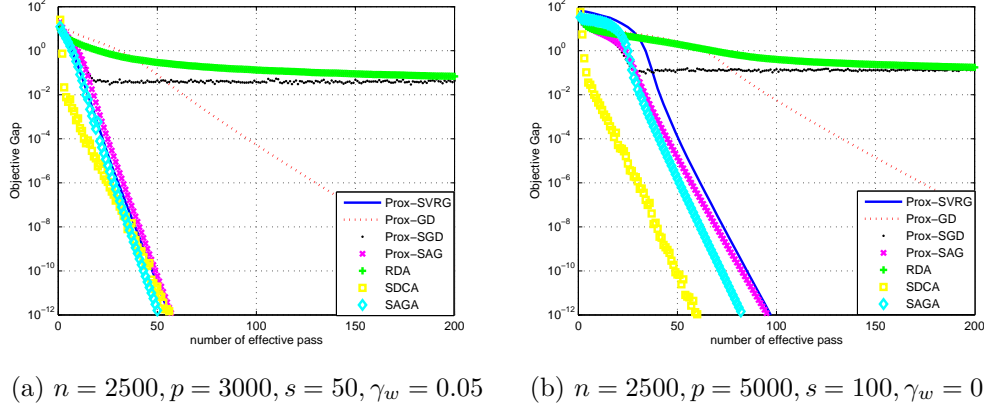


Figure 3: Results on corrected Lasso. The x-axis is the number of pass over the dataset. y-axis is the objective gap $F(w_k) - F(\hat{w})$ with log scale. We try two different settings. In the first figure $n = 2500, p = 3000, s = 50, \gamma_w = 0.05$. In the second figure $n = 2500, p = 5000, s = 100, \gamma_w = 0.1$.

In both setting, we see $\tilde{\kappa} > \tilde{\lambda} + \gamma_\zeta$, thus according to our theory, SDCA converges linearly. In both figures (a) and (b), SDCA, Prox-SVRG, Prox-SAG and Prox-GD converge linearly. SDCA performs better in the second setting. SGD and RDA converge slowly due to the large variance in gradient.

4.1.4 SCAD

The way to generate data is same with Lasso. Here $x_i \in \mathbb{R}^p$ is drawn from normal distribution $N(0, 2I)$ to satisfy the requirement of on $\tilde{\kappa}$, γ_ζ and $\tilde{\lambda}$. We set $\lambda = 0.05$ in the formulation and choose $\tilde{\lambda} = 0.1$ in SDCA. We present the result in Figure 4. We try two different settings on n, p, s, ζ . In both case, SDCA, Prox-SVRG, Prox-SAG, converge linearly with similar performance. Prox-GD also converges linearly but with slower rate. According to our theory, $\tilde{\kappa} \geq 1$ and $\tilde{\lambda} + \mu \leq 0.5$ in both cases, thus SDCA can converge linearly and the simulation results verify our theory.

4.2 Real dataset

4.2.1 Sparse Classification Problem

In this section, we evaluate the performance of the algorithms when solving the logistic regression with ℓ_1 regularization: $\min_w \sum_{i=1}^n \log(1 + \exp(-y_i x_i^T w)) + \lambda \|w\|_1$. We conduct experiments on two real-world data sets, namely, rcv1 Lewis et al. [2004] and sido0 Guyon [2008]. The regularization parameters are set as $\lambda = 2 \cdot 10^{-5}$ in rcv1 and $\lambda = 10^{-4}$ in sido0, as suggested in Xiao and Zhang [2014]. For SDCA, we choose $\tilde{\lambda} = 0.002$ and $\tilde{\lambda} = 0.001$ in these two experiments, respectively.

In Figure 5a and 5b we report the performance of different algorithms. In Figure 5a, Prox-SVRG performs best, and closely followed by SDCA, SAGA, and then Prox-SAG. We observe that Prox-GD converges much slower, albeit in theory it should converge with a linear rate Agarwal et al. [2010], possibly because its contraction factor is close to

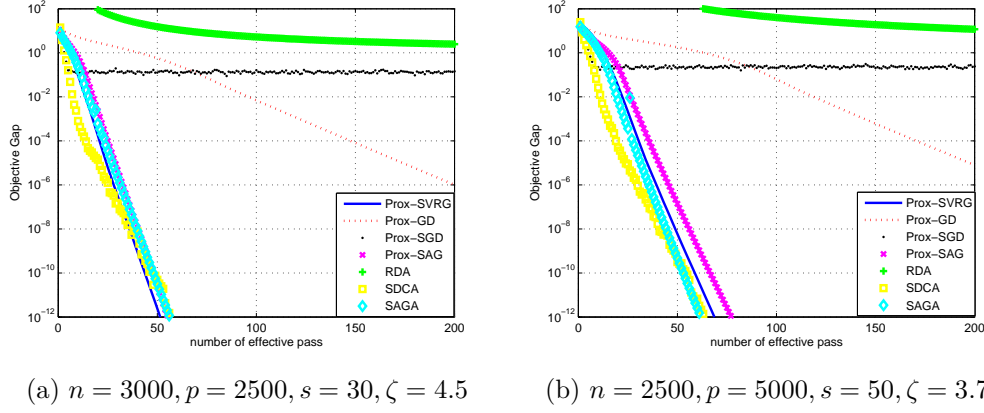


Figure 4: Results on SCAD. The x-axis is the number of pass over the dataset. y-axis is the objective gap $F(w_k) - F(\hat{w})$ with log scale.




Figure 5: The x-axis is the number of pass over the dataset, y-axis is the objective gap in the log-scale.

one. Prox-SGD and RDA converge slowly due to the variance in the stochastic gradient. The objective gaps of them remain significant even after 1000 passes of the whole dataset. In Figure 5b, similarly as before, Prox-SVRG, SAGA, SDCA and Prox-SAG (some part of Prox-SAG overlaps with SDCA) perform well. On this dataset, SAGA performs best followed by SDCA, Prox-SAG and then Prox-SVRG. The performance of Prox-GD is even worse than Prox-SGD. RDA converges the slowest.

4.2.2 Sparse Regression Problem

In Figure 6a, we present the result of Lasso on IJCNN1 dataset ($n = 49990, p = 22$) Prokhorov [2001] with $\lambda=0.02$. It is easy to see, the performance of SDCA is best and then followed by SAGA, Prox-SAG, SVRG and Prox-GD. Prox-SGD and RDA does not work well. We apply the linear regression with SCAD regularization on IJCNN1 dataset Prokhorov [2001] and present the result in Figure 6b. In this dataset, SAGA and SDCA have almost identical performance, then followed by Prox-SVRG, Prox-SAG and Prox-GD. Prox-SGD converge fast at beginning, but has a large optimality gap (10^{-4}). RDA




Figure 6: The x-axis is the number of pass over the dataset, y-axis is the objective gap in the log-scale.

does not work at all. In Figure 6c, we consider a group sparse regression problem on the Boston Housing dataset ($n = 506, p = 13$) Harrison and Rubinfeld [2013]. As suggested in Swirszcz et al. [2009], Xiang et al. [2014], to take into account the non-linear relationship between variables and response, up to third-degree polynomial expansion is applied on each feature. In particular, terms x , x^2 and x^3 are grouped together. We consider group Lasso model on this problem with $\lambda = 0.1$. We choose the setting $m = 2n$ in SVRG and $\tilde{\lambda} = 0.1$ in SDCA. Figure 6c shows the objective gap of various algorithms versus the number of passes over the dataset. Prox-SVRG, SDCA, SAGA and Prox-SAG have almost same performance. Prox-SGD does not converge: the objective gap oscillates between 0.1 and 1. Both the Prox-GD and RDA converge, but with much slower rates.

5 Conclusion and future work

In this paper, we adapt SDCA into a dual-free form to solve non-strongly convex problems and non-convex problems. Under the condition of RSC, we prove that this dual-free SDCA converges with linear rates which covers several important statistical models. From a high level, our results re-confirmed a well-observed phenomenon that statistically easy problems tend to be more computationally friendly. We believe this intriguing fact indicates that there are fundamental relationships between statistics and optimization, and understanding such relationships may shed deep insights.

A Proofs

To begin with, we present a technical Lemma that will be used repeatedly in the following proofs.

Lemma 1. *Suppose function $f(x)$ is convex and L smooth then we have*

$$f(x) - f(y) - \langle \nabla f(y), x - y \rangle \geq \frac{1}{2L} \|\nabla f(x) - \nabla f(y)\|_2^2.$$

Proof. Define $p(x) = f(x) - \langle \nabla f(x_0), x \rangle$. It is obvious that $p(x_0)$ is the optimal value of $p(x)$ due to convexity. Since $f(x)$ is L smooth, so is $p(x)$, and we have

$$p(x_0) \leq p\left(x - \frac{1}{L} \nabla p(x)\right) \leq p(x) - \frac{1}{2L} \|\nabla p(x)\|_2^2.$$

That is

$$f(x_0) - \langle \nabla f(x_0), x_0 \rangle \leq f(x) - \langle \nabla f(x_0), x \rangle - \frac{1}{2L} \|\nabla f(x) - \nabla f(x_0)\|_2^2.$$

Rearrange the terms, we have

$$f(x) - f(x_0) - \langle \nabla f(x_0), x - x_0 \rangle \geq \frac{1}{2L} \|\nabla f(x) - \nabla f(x_0)\|_2^2.$$

□

The following Lemma presents a well-known fact on conjugate function of a strongly convex function. We extract it from Theorem 1 of Nesterov [2005].

Lemma 2. Define $p^*(v) = \max_{w \in \Omega} \langle w, v \rangle - p(w)$, where Ω is a convex compact set, $p(w)$ is a 1-strongly convex function, then $p^*(v)$ is 1-smooth and $\nabla p^*(v) = \bar{w}$ where $\bar{w} = \arg \max_{w \in \Omega} \langle w, v \rangle - p(w)$.

A.1 Proof of results for Convex $F(w)$

In this section we establish the results for convex $F(w)$, namely, Theorem 1. Recall the problem we want to optimize is

$$\min_{w \in \Omega} F(w) := f(w) + \lambda g(w) := \frac{1}{n} \sum_{i=1}^n f_i(w) + \lambda g(w), \quad (7)$$

where Ω is $\{w \in \mathbb{R}^p | g(w) \leq \rho\}$. Remind that instead of directly minimizing the above objective function, we look at the following form.

$$\min_{w \in \Omega} F(w) := \phi(w) + \tilde{\lambda} \tilde{g}(w) = \frac{1}{1+n} \sum_{i=1}^{n+1} \phi_i(w) + \tilde{\lambda} \tilde{g}(w).$$

Also remind that \tilde{L}_i is the smooth parameter of $\phi_i(w)$ and we define $\tilde{L} \triangleq \frac{1}{n+1} \sum_{i=1}^{n+1} \tilde{L}_i$.

Notice the Algorithm 1 keeps the following relation $v^{t-1} = \frac{1}{\lambda(n+1)} \sum_{i=1}^{n+1} a_i^{t-1}$, and thus we have:

$$\begin{aligned} E(\eta_i(\nabla \phi_i(w^{t-1}) + a_i^{t-1})) &= \frac{\eta}{n+1} \sum_{i=1}^{n+1} (\nabla \phi_i(w^{t-1}) + a_i^{t-1}) \\ &= \eta \left(\frac{1}{n+1} \sum_{i=1}^{n+1} \nabla \phi_i(w^{t-1}) + \tilde{\lambda} v^{t-1} \right) \\ &= \eta (\nabla \phi(w^{t-1}) + \tilde{\lambda} v^{t-1}). \end{aligned} \quad (8)$$

Before we start the proof of the main theorem, we present several technical lemmas. The following two lemmas are similar to its batched counterpart in Agarwal et al. [2010].

Lemma 3. Suppose $f(w)$ is convex and $g(w)$ is decomposable with respect to $(\mathcal{A}, \mathcal{B})$, $g(w^*) \leq \rho$, if we choose $\lambda \geq 2g^*(\nabla f(w^*))$, define the error term $\Delta^* = \hat{w} - w^*$, then we have the following condition holds

$$g(\Delta_{\mathcal{B}^\perp}^*) \leq 3g(\Delta_{\mathcal{B}}^*) + 4g(w_{\mathcal{A}^\perp}^*),$$

which implies $g(\Delta^*) \leq g(\Delta_{\mathcal{B}^\perp}^*) + g(\Delta_{\mathcal{B}}^*) \leq 4g(\Delta_{\mathcal{B}}^*) + 4g(w_{\mathcal{A}^\perp}^*)$.

Proof. Using the optimality of \hat{w} , we have

$$f(\hat{w}) + \lambda g(\hat{w}) - f(w^*) - \lambda g(w^*) \leq 0.$$

So we get

$$\lambda g(w^*) - \lambda g(\hat{w}) \geq f(\hat{w}) - f(w^*) \geq \langle \nabla f(w^*), \hat{w} - w^* \rangle \geq -g^*(\nabla f(w^*))g(\Delta^*),$$

where the second inequality holds from the convexity of $f(w)$, and the third one holds by Holder's inequality. Using triangle inequality, we have $g(\Delta^*) \leq g(\Delta_{\mathcal{B}}^*) + g(\Delta_{\mathcal{B}^\perp}^*)$, which leads to

$$\lambda g(w^*) - \lambda g(\hat{w}) \geq -g^*(\nabla f(w^*))(g(\Delta_{\mathcal{B}}^*) + g(\Delta_{\mathcal{B}^\perp}^*)). \quad (9)$$

Notice

$$\hat{w} = w^* + \Delta^* = w_{\mathcal{A}}^* + w_{\mathcal{A}^\perp}^* + \Delta_{\mathcal{B}}^* + \Delta_{\mathcal{B}^\perp}^*.$$

Now we obtain

$$\begin{aligned} g(\hat{w}) - g(w^*) &\stackrel{(a)}{\geq} g(w_{\mathcal{A}}^* + \Delta_{\mathcal{B}^\perp}^*) - g(w_{\mathcal{A}^\perp}^*) - g(\Delta_{\mathcal{B}}^*) - g(w^*) \\ &\stackrel{(b)}{=} g(w_{\mathcal{A}}^*) + g(\Delta_{\mathcal{B}^\perp}^*) - g(w_{\mathcal{A}^\perp}^*) - g(\Delta_{\mathcal{B}}^*) - g(w^*) \\ &\stackrel{(c)}{\geq} g(w_{\mathcal{A}}^*) + g(\Delta_{\mathcal{B}^\perp}^*) - g(w_{\mathcal{A}^\perp}^*) - g(\Delta_{\mathcal{B}}^*) - g(w_{\mathcal{A}}^*) - g(w_{\mathcal{A}^\perp}^*) \\ &\geq g(\Delta_{\mathcal{B}^\perp}^*) - 2g(w_{\mathcal{A}^\perp}^*) - g(\Delta_{\mathcal{B}}^*), \end{aligned} \quad (10)$$

where (a) and (c) hold from the triangle inequality, and (b) uses the decomposability of $g(\cdot)$.

Substitute the above to (9), and use the assumption that $\lambda \geq 2g^*(\nabla f(w^*))$, we obtain

$$-\frac{\lambda}{2}(g(\Delta_{\mathcal{B}}^*) + g(\Delta_{\mathcal{B}^\perp}^*)) + \lambda(g(\Delta_{\mathcal{B}^\perp}^*) - 2g(w_{\mathcal{A}^\perp}^*) - g(\Delta_{\mathcal{B}}^*)) \leq 0,$$

which implies

$$g(\Delta_{\mathcal{B}^\perp}^*) \leq 3g(\Delta_{\mathcal{B}}^*) + 4g(w_{\mathcal{A}^\perp}^*).$$

□

Lemma 4. Suppose $f(w)$ is convex and $g(w)$ is decomposable with respect to $(\mathcal{A}, \mathcal{B})$. If we choose $\lambda \geq 2g^*(\nabla f(w^*))$, and suppose there exist a given tolerance ξ and T such that $F(w^t) - F(\hat{w}) \leq \xi$ for all $t > T$, then for the error term $\Delta^t = w^t - w^*$ we have

$$g(\Delta_{\mathcal{B}^\perp}^t) \leq 3g(\Delta_{\mathcal{B}}^t) + 4g(w_{\mathcal{A}^\perp}^*) + 2\min\{\frac{\xi}{\lambda}, \rho\},$$

which implies

$$g(\Delta^t) \leq 4g(\Delta_{\mathcal{B}}^t) + 4g(w_{\mathcal{A}^\perp}^*) + 2\min\{\frac{\xi}{\lambda}, \rho\}.$$

Proof. First notice $F(w^t) - F(w^*) \leq \xi$ holds by assumption since $F(w^*) \geq F(\hat{w})$. So we have

$$f(w^t) + \lambda g(w^t) - f(w^*) - \lambda g(w^*) \leq \xi.$$

Follow the same steps as those in the proof of Lemma 3, we have

$$g(\Delta_{\mathcal{B}^\perp}^t) \leq 3g(\Delta_{\mathcal{B}}^t) + 4g(w_{\mathcal{A}^\perp}^*) + 2\frac{\xi}{\lambda}.$$

Using the fact that w^* and w^t both belong to $\Omega = \{w \in \mathbb{R}^p | g(w) \leq \rho\}$, we have $g(\Delta^t) \leq g(w^*) + g(w^t) \leq 2\rho$. This leads to

$$g(\Delta_{\mathcal{B}^\perp}^t) \leq g(\Delta_{\mathcal{B}}^t) + 2\rho,$$

by triangle inequality $g(\Delta_{\mathcal{B}^\perp}^t) \leq g(\Delta_{\mathcal{B}}^t) + g(\Delta^t)$. Combine this with the above result we have

$$g(\Delta_{\mathcal{B}^\perp}^t) \leq 3g(\Delta_{\mathcal{B}}^t) + 4g(w_{\mathcal{A}^\perp}^*) + 2\min\{\frac{\xi}{\lambda}, \rho\}.$$

The second statement follows immediately from $g(\Delta^t) \leq g(\Delta_{\mathcal{B}}^t) + g(\Delta_{\mathcal{B}^\perp}^t)$. \square

Lemma 5. *Under the same assumption of Lemma 4, we have*

$$F(w^t) - F(\hat{w}) \geq \left(\frac{\kappa}{2} - 32\tau\Psi^2(\mathcal{B})\right) \|\hat{\Delta}^t\|_2^2 - \epsilon^2(\Delta^*, \mathcal{A}, \mathcal{B}), \quad (11)$$

and

$$\begin{aligned} & \phi(\hat{w}) - \phi(w^t) - \langle \hat{w} - w^t, \nabla \phi(w^t) \rangle \\ & \geq \left[\left(\frac{\kappa - \tilde{\lambda}}{2} - 32\tau\Psi^2(\mathcal{B}) \right) \|\hat{\Delta}^t\|_2^2 - \epsilon^2(\Delta^*, \mathcal{A}, \mathcal{B}) \right], \end{aligned} \quad (12)$$

where $\hat{\Delta}^t = w^t - \hat{w}$, $\epsilon^2(\Delta^*, \mathcal{A}, \mathcal{B}) = 2\tau(\delta_{stat} + \delta)^2$, $\delta = 2\min\{\frac{\xi}{\lambda}, \rho\}$, and $\delta_{stat} = 8\Psi(\mathcal{B})\|\Delta^*\|_2 + 8g(w_{\mathcal{A}^\perp}^*)$.

Proof. We begin the proof by establishing a simple fact on $\hat{\Delta}^t = w^t - \hat{w}$. We adapt the argument in Lemma 4 (which is on Δ^t) to $\hat{\Delta}^t$:

$$\begin{aligned} g(\hat{\Delta}^t) & \leq g(\Delta^t) + g(\Delta^*) \\ & \leq 4g(\Delta_{\mathcal{B}}^t) + 4g(w_{\mathcal{A}^\perp}^*) + 2\min\{\frac{\xi}{\lambda}, \rho\} + 4g(\Delta_{\mathcal{B}}^*) + 4g(w_{\mathcal{A}^\perp}^*) \\ & \leq 4\Psi(\mathcal{B})\|\Delta^t\|_2 + 4\Psi(\mathcal{B})\|\Delta^*\|_2 + 8g(w_{\mathcal{A}^\perp}^*) + 2\min\{\frac{\xi}{\lambda}, \rho\}, \end{aligned} \quad (13)$$

where the first inequality holds from the triangle inequality, the second inequality uses Lemma 3 and 4, the third holds because of the definition of subspace compatibility.

We know

$$f(w^t) - f(\hat{w}) - \langle \nabla f(\hat{w}), \hat{\Delta}^t \rangle \geq \frac{\kappa}{2} \|\hat{\Delta}^t\|_2^2 - \tau g^2(\hat{\Delta}^t)$$

which implies $F(w^t) - F(\hat{w}) \geq \frac{\kappa}{2} \|\hat{\Delta}^t\|_2^2 - \tau g^2(\hat{\Delta}^t)$, since \hat{w} is the optimal solution to the problem 7 and $g(w)$ is convex. Notice that

$$\begin{aligned} g(\hat{\Delta}^t) & \leq 4\Psi(\mathcal{B})\|\Delta^t\|_2 + 4\Psi(\mathcal{B})\|\Delta^*\|_2 + 8g(w_{\mathcal{A}^\perp}^*) + 2\min\{\frac{\xi}{\lambda}, \rho\} \\ & \leq 4\Psi(\mathcal{B})\|\hat{\Delta}^t\|_2 + 8\Psi(\mathcal{B})\|\Delta^*\|_2 + 8g(w_{\mathcal{A}^\perp}^*) + 2\min\{\frac{\xi}{\lambda}, \rho\}, \end{aligned} \quad (14)$$

where the second inequality uses the triangle inequality. Using the inequality $(a + b)^2 \leq 2a^2 + 2b^2$, we can upper bound $g^2(\hat{\Delta}^t)$.

$$g^2(\hat{\Delta}^t) \leq 32\Psi^2(\mathcal{B})\|\hat{\Delta}^t\|_2^2 + 2[8\Psi(\mathcal{B})\|\Delta^*\|_2 + 8g(w_{\mathcal{A}^\perp}^*) + 2\min\{\frac{\xi}{\lambda}, \rho\}]^2. \quad (15)$$

We now use above result to rewrite the RSC condition.

Substitute this upper bound in the RSC, we have

$$F(w^t) - F(\hat{w}) \geq \left(\frac{\kappa}{2} - 32\tau\Psi^2(\mathcal{B}) \right) \|\hat{\Delta}^t\|_2^2 - 2\tau[8\Psi(\mathcal{B})\|\Delta^*\|_2 + 8g(w_{\mathcal{A}^\perp}^*) + 2\min\{\frac{\xi}{\lambda}, \rho\}]^2.$$

Notice by $\delta = 2\min\{\frac{\xi}{\lambda}, \rho\}$, $\delta_{stat} = 8\Psi(\mathcal{B})\|\Delta^*\|_2 + 8g(w_{\mathcal{A}^\perp}^*)$, and $\epsilon^2(\Delta^*, \mathcal{A}, \mathcal{B}) = 2\tau(\delta_{stat} + \delta)^2$, we obtain

$$\epsilon^2(\Delta^*, \mathcal{A}, \mathcal{B}) = 2\tau(8\Psi(\mathcal{B})\|\Delta^*\|_2 + 8g(w_{\mathcal{A}^\perp}^*) + 2\min\{\frac{\xi}{\lambda}, \rho\})^2.$$

We thus conclude

$$F(w^t) - F(\hat{w}) \geq \left(\frac{\kappa}{2} - 32\tau\Psi^2(\mathcal{B}) \right) \|\hat{\Delta}^t\|_2^2 - \epsilon^2(\Delta^*, \mathcal{A}, \mathcal{B}). \quad (16)$$

Recall $\phi(w) = f(w) - \frac{\tilde{\lambda}}{2}\|w\|_2^2$, and hence we have

$$\begin{aligned} & \phi(\hat{w}) - \phi(w^t) - \langle \hat{w} - w^t, \nabla\phi(w^t) \rangle \\ &= f(\hat{w}) - f(w^t) - \langle \nabla f(w^t), \hat{w} - w^t \rangle - \frac{\tilde{\lambda}}{2}\|\hat{w} - w^t\|_2^2 \\ &\geq \frac{\kappa}{2}\|\hat{\Delta}^t\|_2^2 - \tau g^2(\hat{\Delta}^t) - \frac{\tilde{\lambda}}{2}\|\hat{\Delta}^t\|_2^2, \end{aligned} \quad (17)$$

where the inequality is due to the RSC condition. Now we plug in the upper bound of $g^2(\hat{\Delta}^t)$ in Equation (15), and arrange the terms to establish Equation (12). \square

Proof of Theorem on convex $F(w)$, i.e., Theorem 1. Recall we define two potentials

$$\begin{aligned} A_t &= \sum_{j=1}^{n+1} \frac{1}{q_i} \|a_j^t - \hat{a}_j\|_2^2, \\ B_t &= 2(\tilde{g}^*(v^t) - \langle \nabla\tilde{g}^*(\hat{v}), v^t - \hat{v} \rangle - \tilde{g}^*(\hat{v})). \end{aligned}$$

Notice using Theorem 1 in Nesterov [2005], we know $\tilde{g}^*(v)$ is 1-smooth and $w = \nabla\tilde{g}^*(v)$.

We remark that the potential A_t is defined the same as in Shalev-Shwartz [2016] while we define B_t differently to solve the problem with general regularization $g(w)$. When $\tilde{g}(w) = \frac{1}{2}\|w\|_2^2$, B_t reduce to $\|v^t - \hat{v}\|_2^2 = \|w^t - \hat{w}\|_2^2$, which is same as in Shalev-Shwartz [2016].

Step 1. The first step is to lower bound $A_{t-1} - A_t$, in particular, to establish that

$$\mathbb{E}[A_{t-1} - A_t] = \eta\tilde{\lambda} \left(A_{t-1} + \sum_{i=1}^{n+1} \frac{1}{q_i} (-\|u_i - \hat{a}_i\|_2^2 + (1 - \beta_i)\|m_i\|_2^2) \right). \quad (18)$$

This step is indeed same as Shalev-Shwartz [2016], which we present for the completeness. Define $u_i = -\nabla\phi_i(w^{t-1})$, $\beta_i = \eta_i\tilde{\lambda}(n+1)$ and $m_i = -u_i + a_i^{t-1}$ for notational simplicity. Suppose coordinate i is picked up at time t , then we have

$$\begin{aligned} A_{t-1} - A_t &= -\frac{1}{q_i} \|a_i^t - \hat{a}_i\|_2^2 + \frac{1}{q_i} \|a_i^{t-1} - \hat{a}_i\|_2^2 \\ &= -\frac{1}{q_i} \|(1 - \beta_i)(a_i^{t-1} - \hat{a}_i) + \beta_i(u_i - \hat{a}_i)\|_2^2 + \frac{1}{q_i} \|a_i^{t-1} - \hat{a}_i\|_2^2 \\ &\geq -\frac{1}{q_i} ((1 - \beta_i)\|a_i^{t-1} - \hat{a}_i\|_2^2 + \beta_i\|u_i - \hat{a}_i\|_2^2 - \beta_i(1 - \beta_i)\|a_i^{t-1} - u_i\|_2^2) \\ &\quad + \frac{1}{q_i} \|a_i^{t-1} - \hat{a}_i\|_2^2 \\ &= \frac{\beta_i}{q_i} (\|a_i^{t-1} - \hat{a}_i\|_2^2 - \|u_i - \hat{a}_i\|_2^2 + (1 - \beta_i)\|m_i\|_2^2). \end{aligned} \quad (19)$$

Taking expectation on both sides with respect to the random sampling of i at time step t , we established Equation (18).

Step 2. We now look at the evolution of B_t . In particular, we will prove that

$$\mathbb{E}(B_{t-1} - B_t) = 2\langle w^{t-1} - \hat{w}, \eta(\nabla\phi(w^{t-1}) + \tilde{\lambda}v^{t-1}) \rangle - \frac{\eta^2}{(n+1)^2} \sum_{i=1}^{n+1} \frac{1}{q_i} \|m_i\|_2^2. \quad (20)$$

To this end, notice that

$$\begin{aligned} & \frac{1}{2}(B_{t-1} - B_t) \\ &= \tilde{g}^*(v^{t-1}) - \langle \nabla\tilde{g}^*(\hat{v}), v^{t-1} - \hat{v} \rangle - \tilde{g}^*(\hat{v}) - (\tilde{g}^*(v^t) - \langle \nabla\tilde{g}^*(\hat{v}), v^t - \hat{v} \rangle - \tilde{g}^*(\hat{v})) \\ &= \tilde{g}^*(v^{t-1}) - \tilde{g}^*(v^t) - \langle \nabla\tilde{g}^*(\hat{v}), v^{t-1} - v^t \rangle \\ &\geq \langle \nabla\tilde{g}^*(v^{t-1}), v^{t-1} - v^t \rangle - \frac{1}{2}\|v^{t-1} - v^t\|_2^2 - \langle \nabla\tilde{g}^*(\hat{v}), v^{t-1} - v^t \rangle \\ &= \langle \nabla\tilde{g}^*(v^{t-1}) - \nabla\tilde{g}^*(\hat{v}), v^{t-1} - v^t \rangle - \frac{1}{2}\|v^{t-1} - v^t\|_2^2 \\ &= \langle w^{t-1} - \hat{w}, v^{t-1} - v^t \rangle - \frac{1}{2}\|v^{t-1} - v^t\|_2^2, \end{aligned} \quad (21)$$

where the inequality holds because

$$\tilde{g}^*(v^t) - \tilde{g}^*(v^{t-1}) \leq \langle \nabla\tilde{g}^*(v^{t-1}), v^t - v^{t-1} \rangle + \frac{1}{2}\|v^t - v^{t-1}\|_2^2,$$

which results from $\tilde{g}^*(v)$ being 1-smooth; and the last equality holds from the fact that $\nabla\tilde{g}^*(v^{t-1}) = w^{t-1}$ using Lemma 2.

Take expectation on both sides of Equation (21), we established Equation (20), where we use result in Equation (8).

Step 3. We define a new potential $C_t = c_a A_t + c_b B_t$, and prove in this step that

$$\begin{aligned} \mathbb{E}(C_{t-1} - C_t) &\geq c_a \eta \tilde{\lambda} A_{t-1} - c_a \eta \tilde{\lambda} \sum_{i=1}^{n+1} \frac{1}{q_i} \|u_i - \hat{a}_i\|_2^2 + c_b \eta \tilde{\lambda} B_{t-1} \\ &\quad + 2c_b \eta (F(w^{t-1}) - F(\hat{w})) + 2c_b \eta (\phi(\hat{w}) - \phi(w^{t-1}) - \langle \hat{w} - w^{t-1}, \nabla\phi(w^{t-1}) \rangle). \end{aligned} \quad (22)$$

From the definition of C_t , and using Equation (18) and (20), we have

$$\begin{aligned} \mathbb{E}(C_{t-1} - C_t) &\geq c_a \eta \tilde{\lambda} A_{t-1} - c_a \eta \tilde{\lambda} \sum_{i=1}^{n+1} \frac{1}{q_i} \|u_i - \hat{a}_i\|_2^2 + 2c_b \eta \langle w^{t-1} - \hat{w}, \nabla\phi(w^{t-1}) + \tilde{\lambda}v^{t-1} \rangle \\ &\quad + \sum_{i=1}^{n+1} \frac{1}{q_i} \|m_i\|_2^2 (c_a \eta \tilde{\lambda} (1 - \beta_i) - \frac{c_b \eta^2}{(n+1)^2}). \end{aligned} \quad (23)$$

We next show that

$$\langle w^{t-1} - \hat{w}, \nabla\phi(w^{t-1}) + \tilde{\lambda}v^{t-1} \rangle - \frac{\tilde{\lambda}}{2} B_{t-1} - (F(w^{t-1}) - F(\hat{w})) = \phi(\hat{w}) - \phi(w^{t-1}) - \langle \hat{w} - w^{t-1}, \nabla\phi(w^{t-1}) \rangle.$$

This holds by directly verifying as follows:

$$\begin{aligned}
& \langle w^{t-1} - \hat{w}, \nabla \phi(w^{t-1}) + \tilde{\lambda} v^{t-1} \rangle - \frac{\tilde{\lambda}}{2} B_{t-1} - (F(w^{t-1}) - F(\hat{w})) \\
&= \langle w^{t-1} - \hat{w}, \nabla \phi(w^{t-1}) + \tilde{\lambda} v^{t-1} \rangle - \tilde{\lambda} (\tilde{g}^*(v^{t-1}) - \langle \nabla \tilde{g}^*(\hat{v}), v^{t-1} - \hat{v} \rangle - \tilde{g}^*(\hat{v})) - F(w^{t-1}) + F(\hat{w}) \\
&= \langle w^{t-1}, \tilde{\lambda} v^{t-1} \rangle - \tilde{\lambda} \tilde{g}^*(v^{t-1}) - \langle \hat{w}, \tilde{\lambda} \hat{v} \rangle + \tilde{\lambda} \tilde{g}^*(\hat{v}) + \langle \hat{w}, -\nabla \phi(w^{t-1}) \rangle - F(w^{t-1}) + F(\hat{w}) \\
&\quad + \langle w^{t-1}, \nabla \phi(w^{t-1}) \rangle \\
&= \tilde{\lambda} \tilde{g}(w^{t-1}) - \tilde{\lambda} \tilde{g}(\hat{w}) - F(w^{t-1}) + F(\hat{w}) + \langle w^{t-1} - \hat{w}, \nabla \phi(w^{t-1}) \rangle \\
&= \tilde{\lambda} \tilde{g}(w^{t-1}) - \tilde{\lambda} \tilde{g}(\hat{w}) - \tilde{\lambda} \tilde{g}(w^{t-1}) - \phi(w^{t-1}) + \tilde{\lambda} \tilde{g}(\hat{w}) + \phi(\hat{w}) + \langle w^{t-1} - \hat{w}, \nabla \phi(w^{t-1}) \rangle \\
&= \phi(\hat{w}) - \phi(w^{t-1}) - \langle \hat{w} - w^{t-1}, \nabla \phi(w^{t-1}) \rangle,
\end{aligned} \tag{24}$$

where the second equality uses the fact that $\hat{w} = \nabla \tilde{g}^*(\hat{v})$, and the third equality holds using the definition of $\tilde{g}^*(v)$. Thus, substituting the equation into (23), we get

$$\begin{aligned}
\mathbb{E}(C_{t-1} - C_t) &\geq c_a \eta \tilde{\lambda} A_{t-1} - c_a \eta \tilde{\lambda} \sum_{i=1}^{n+1} \frac{1}{q_i} \|u_i - \hat{a}_i\|_2^2 + c_b \eta \tilde{\lambda} B_{t-1} + 2c_b \eta (F(w^{t-1}) - F(\hat{w})) \\
&\quad + \sum_{i=1}^{n+1} \frac{1}{q_i} \|m_i\|_2^2 (c_a \eta \tilde{\lambda} (1 - \beta_i) - \frac{c_b \eta^2}{(n+1)^2}) + 2c_b \eta [\phi(\hat{w}) - \phi(w^{t-1}) - \langle \hat{w} - w^{t-1}, \nabla \phi(w^{t-1}) \rangle].
\end{aligned} \tag{25}$$

We can choose $\eta \leq \frac{q_i}{2\tilde{\lambda}}$ and $\frac{c_b}{c_a} = \frac{\tilde{\lambda}(n+1)^2}{2\eta}$ so that $\beta_i \leq 1/2$, and the term $\sum_{i=1}^{n+1} \frac{1}{q_i} \|m_i\|_2^2 (c_a \eta \tilde{\lambda} (1 - \beta_i) - \frac{c_b \eta^2}{(n+1)^2})$ is non-negative. Since $q_i \geq \frac{1}{2(n+1)}$ for every i , we can choose $\eta \leq \frac{1}{4\tilde{\lambda}(n+1)}$. Using these condition, we established (22).

Step 4 We now bound $\sum_{i=1}^{n+1} \frac{1}{q_i} \|u_i - \hat{a}_i\|_2^2$. Notice that for $i = 1, \dots, n$, since ϕ_i is convex, we can apply Lemma 1, and the fact that $\xi + \nabla f(\hat{w}) = 0$, for $\xi \in \lambda \partial g(\hat{w})$.

$$\begin{aligned}
\sum_{i=1}^n \frac{1}{q_i} \|u_i - \hat{a}_i\|_2^2 &= \sum_{i=1}^n \frac{1}{q_i} \|\nabla \phi_i(w^{t-1}) - \nabla \phi_i(\hat{w})\|_2^2 \\
&\leq \left(2 \max \frac{\tilde{L}_i}{q_i} \right) \sum_{i=1}^n (\phi_i(w^{t-1}) - \phi_i(\hat{w}) - \langle \nabla \phi_i(\hat{w}), w^{t-1} - \hat{w} \rangle) \\
&= \frac{n+1}{n} \left(2 \max \frac{\tilde{L}_i}{q_i} \right) \sum_{i=1}^n (f_i(w^{t-1}) - f_i(\hat{w}) - \langle \nabla f_i(\hat{w}), w^{t-1} - \hat{w} \rangle) \tag{26} \\
&\leq \frac{n+1}{n} \left(2 \max \frac{\tilde{L}_i}{q_i} \right) n (f(w^{t-1}) - f(\hat{w}) + \lambda g(w^{t-1}) - \lambda g(\hat{w})) \\
&\leq \left(2 \max \frac{\tilde{L}_i}{q_i} \right) (n+1) (F(w^{t-1}) - F(\hat{w})).
\end{aligned}$$

As for $i = n+1$, we have

$$\begin{aligned}
\frac{1}{q_{n+1}} \|\nabla \phi_{n+1}(w^{t-1}) - \nabla \phi_{n+1}(\hat{w})\|_2^2 &\leq \frac{\tilde{\lambda}^2 (n+1)^2}{q_{n+1}} \|w^{t-1} - \hat{w}\|_2^2 \\
&= 2(n+1) \frac{\tilde{L}_{n+1}}{q_{n+1}} \frac{\tilde{\lambda}}{2} \|w^{t-1} - \hat{w}\|_2^2.
\end{aligned} \tag{27}$$

Step 5 We now analyze the progress of the potential by induction. We need to relate $\frac{\tilde{\lambda}}{2}\|w^{t-1} - \hat{w}\|_2^2$ to $F(w^{t-1}) - F(\hat{w})$. In high level, we divide the time steps $t = 1, 2, \dots$ into several epochs, i.e., $([T_0, T_1), [T_1, T_2), \dots)$. At the end of each epoch j , we prove that C_t decreases with linear rate until the optimality gap $F(w^t) - F(\hat{w})$ decrease to some tolerance ξ_j . We then prove that $(\xi_1, \xi_2, \xi_3, \dots)$ is a decreasing sequence and finish the proof.

Assuming that time step $t - 1$ is in the epoch j , we use Equation (16) in Lemma 5 and the fact that $\tilde{\lambda} \leq \tilde{\kappa}$ to obtain

$$\frac{1}{q_{n+1}}\|\nabla\phi_{n+1}(w^{t-1}) - \nabla\phi_{n+1}(\hat{w})\|_2^2 \leq 2(n+1)\frac{\tilde{L}_{n+1}}{q_{n+1}}(F(w^{t-1}) - F(\hat{w}) + \epsilon_j^2(\Delta^*, \mathcal{A}, \mathcal{B})).$$

Combining the above two results on $\phi_i(\cdot)$ together, we have

$$\begin{aligned} \sum_{i=1}^{n+1} \frac{1}{q_i} \|u_i - \hat{a}_i\|_2^2 &\leq 4(n+1) \left(\max_{i \in \{1, \dots, n+1\}} \frac{\tilde{L}_i}{q_i} \right) (F(w^{t-1}) - F(\hat{w}) + \frac{1}{2}\epsilon_j^2(\Delta^*, \mathcal{A}, \mathcal{B})) \\ &\leq 8(n+1)^2 \tilde{L} (F(w^{t-1}) - F(\hat{w})) + 4(n+1)^2 \tilde{L} \epsilon_j^2(\Delta^*, \mathcal{A}, \mathcal{B}), \end{aligned} \quad (28)$$

where we use the fact $\frac{\tilde{L}_i}{q_i} = 2(n+1)\tilde{L}\frac{\tilde{L}_i}{L_i + L} \leq 2(n+1)\tilde{L}$.

Replace the corresponding term in equation (22), we have

$$\begin{aligned} \mathbb{E}(C_{t-1} - C_t) &\geq c_a \eta \tilde{\lambda} A_{t-1} - 8c_a \eta \tilde{\lambda} (n+1)^2 \tilde{L} (F(w^{t-1}) - F(\hat{w})) - 4c_a \eta \tilde{\lambda} (n+1)^2 \tilde{L} \epsilon_j^2(\Delta^*, \mathcal{A}, \mathcal{B}) \\ &\quad + c_b \eta \tilde{\lambda} B_{t-1} + 2c_b \eta (F(w^{t-1}) - F(\hat{w})) + 2c_b \eta (\phi(\hat{w}) - \phi(w^{t-1}) - \langle \hat{w} - w^{t-1}, \nabla \phi(w^{t-1}) \rangle) \\ &\geq \eta \tilde{\lambda} C_{t-1} + \eta (2c_b - 8c_a \tilde{\lambda} (n+1)^2 \tilde{L}) (F(w^{t-1}) - F(\hat{w})) - 4c_a \eta \tilde{\lambda} (n+1)^2 \tilde{L} \epsilon_j^2(\Delta^*, \mathcal{A}, \mathcal{B}) \\ &\quad + (c_b \eta (\tilde{\kappa} - \tilde{\lambda}) \|w^t - \hat{w}\|_2^2 - 2c_b \eta \epsilon_j^2(\Delta^*, \mathcal{A}, \mathcal{B})), \end{aligned} \quad (29)$$

where the second inequality is due to Lemma 5. We choose $2c_b = 16c_a \tilde{\lambda} (n+1)^2 \tilde{L}$, and use the fact that $\frac{c_b}{c_a} = \frac{\tilde{\lambda}(n+1)^2}{2\eta}$, we have $\eta = \frac{1}{16\tilde{L}}$ and

$$\mathbb{E}(C_{t-1} - C_t) \geq \eta \tilde{\lambda} C_{t-1} + \eta c_b (F(w^{t-1}) - F(\hat{w})) - 3c_b \eta \epsilon_j^2(\Delta^*, \mathcal{A}, \mathcal{B}),$$

where we use the assumption that $\tilde{\kappa} \geq \tilde{\lambda}$ and the assumption $n > 4$.

Remind that in epoch j , we have

$$\begin{aligned} \epsilon_j^2(\Delta^*, \mathcal{A}, \mathcal{B}) &= 2\tau(\delta_{stat} + \delta_{j-1})^2, \delta_{j-1} = 2\min\{\frac{\xi_{j-1}}{\lambda}, \rho\}, \\ \text{and } \delta_{stat} &= 8\Psi(\mathcal{B})\|\Delta^*\| + 8g(w_{\mathcal{A}^\perp}^*). \end{aligned}$$

The epoch is determined by the comparison between $F(w^t) - F(\hat{w})$ and $3\epsilon_j^2(\Delta^*, \mathcal{A}, \mathcal{B})$

In the first epoch, we choose $\delta_0 = 2\rho$. Thus $\epsilon_1^2(\Delta^*, \mathcal{A}, \mathcal{B}) = 2\tau(\delta_{stat} + 2\rho)^2$. We choose T_1 such that

$$F(w^{T_1-1}) - F(\hat{w}) \geq 3\epsilon_1^2(\Delta^*, \mathcal{A}, \mathcal{B}) \quad \text{and} \quad F(w^{T_1}) - F(\hat{w}) \leq 3\epsilon_1^2(\Delta^*, \mathcal{A}, \mathcal{B}).$$

If no such T_1 exists, that means $F(w^t) - F(\hat{w}) \geq 3\epsilon_1^2(\Delta^*, \mathcal{A}, \mathcal{B})$ always holds and we have $\mathbb{E}(C_t) \leq (1 - \eta\tilde{\lambda})C_{t-1}$ for all t and get the geometric convergence rate. It is a contradiction with $F(w^t) - F(\hat{w}) \geq 3\epsilon_1^2(\Delta^*, \mathcal{A}, \mathcal{B})$.

Now we know $F(w^{T_1}) - F(\hat{w}) \leq 3\epsilon_1^2(\Delta^*, \mathcal{A}, \mathcal{B})$, and hence we choose $\xi_1 = 6\tau(\delta_{stat} + \delta_0)^2$.

In the second epoch we use the same argument:

$$\epsilon_2^2(\Delta^*, \mathcal{A}, \mathcal{B}) = 2\tau_\sigma(\delta_{stat} + \delta_1)^2 \quad \text{where} \quad \delta_1 = 2 \min\{\frac{\xi_1}{\lambda}, \rho\}.$$

We choose T_2 such that

$$F(w^{T_2-1}) - F(\hat{w}) \geq 3\epsilon_2^2(\Delta^*, \mathcal{A}, \mathcal{B}) \quad \text{and} \quad F(w^{T_2}) - F(\hat{w}) \leq 3\epsilon_2^2(\Delta^*, \mathcal{A}, \mathcal{B}).$$

Then we choose $\xi_2 = 6\tau(\delta_{stat} + \delta_1)^2$.

Similarly in epoch j , we choose T_j such that

$$F(w^{T_j-1}) - F(\hat{w}) \geq 3\epsilon_j^2(\Delta^*, \mathcal{A}, \mathcal{B}) \quad \text{and} \quad F(w^{T_j}) - F(\hat{w}) \leq 3\epsilon_j^2(\Delta^*, \mathcal{A}, \mathcal{B}),$$

and $\xi_j = 6\tau(\delta_{stat} + \delta_{j-1})^2$.

In this way, we arrive at recursive equalities of the tolerance $\{\xi_j\}_{j=1}^\infty$ where

$$\xi_j = 6\tau_\sigma(\delta_{stat} + \delta_{j-1})^2 \quad \text{and} \quad \delta_{j-1} = 2 \min\{\frac{\xi_{j-1}}{\lambda}, \rho\}.$$

We claim that following holds, until $\delta_j = \delta_{stat}$.

$$\begin{aligned} (I) \quad & \xi_{k+1} \leq \xi_k / (4^{2^{k-1}}) \\ \text{and } (II) \quad & \frac{\xi_{k+1}}{\lambda} \leq \frac{\rho}{4^{2^k}} \quad \text{for } k = 1, 2, 3, \dots \end{aligned} \tag{30}$$

The proof of Equation (30) is same with Equation (60) in Agarwal et al. [2010], which we present here for completeness.

We assume $\delta_0 \geq \delta_{stat}$ (otherwise the statement is true trivially), so $\xi_1 \leq 96\tau\rho^2$. We make the assumption that $\lambda \geq 384\tau\rho$, so $\frac{\xi_1}{\lambda} \leq \frac{\rho}{4}$ and $\xi_1 \leq \xi_0$.

In the second epoch we have

$$\xi_2 \stackrel{(1)}{\leq} 12\tau(\delta_{stat}^2 + \delta_1^2) \leq 24\tau\delta_1^2 \leq \frac{96\tau\xi_1^2}{\lambda^2} \stackrel{(2)}{\leq} \frac{96\tau\xi_1}{4\lambda} \stackrel{(3)}{\leq} \frac{\xi_1}{4},$$

where (1) holds from the fact that $(a+b)^2 \leq 2a^2 + 2b^2$, (2) holds using $\frac{\xi_1}{\lambda} \leq \frac{\rho}{4}$, (3) uses the assumption on λ . Thus,

$$\frac{\xi_2}{\lambda} \leq \frac{\xi_1}{4\lambda} \leq \frac{\rho}{16}.$$

In $i+1$ th step, with similar argument, and by induction assumption we have

$$\xi_{j+1} \leq \frac{96\tau\xi_j^2}{\lambda^2} \leq \frac{96\tau\xi_j}{4^{2^j}\lambda} \leq \frac{\xi_j}{4^{2^{j-1}}}$$

and

$$\frac{\xi_{j+1}}{\lambda} \leq \frac{\xi_j}{4^{2^{j-1}}\lambda} \leq \frac{\rho}{4^{2^j}}.$$

Thus we know ξ_j is a decreasing sequence, and $\mathbb{E}(C_t) \leq (1 - \eta\tilde{\lambda})C_{t-1}$ holds until $F(w^t) - F(\hat{w}) \leq 6\tau(2\delta_{stat})^2$, where η is set to satisfy $\eta \leq \min(\frac{1}{16\tilde{L}}, \frac{1}{4\tilde{\lambda}(n+1)})$.

Observe that $(n+1)\tilde{L} = \frac{n+1}{n}(\sum_{i=1}^n L_i) + \lambda(n+1) = (n+1)(\lambda + \frac{1}{n}\sum_{i=1}^n L_i)$, we establish the theorem. \square

A.2 Strongly convex $F(w)$

Proof of Proposition 1. If $g(w)$ is 1 strongly convex, we can directly apply the algorithm on Equation (7). In this case, RSC is not needed and the proof is significantly simpler. As the main road map of the proof is similar, we just mention the difference in the following. We define the potential,

$$A_t = \sum_{j=1}^n \frac{1}{q_i} \|a_j^t - \hat{a}_j\|_2^2,$$

$$B_t = 2(g^*(v^t) - \langle \nabla g^*(\hat{v}), v^t - \hat{v} \rangle - g^*(\hat{v})).$$

Notice in this setup, the potential B_t is defined on g rather than \tilde{g} .

The evolution of A_t follows from the same analysis of Step 1 in the proof of Theorem 1, except that we replace $\tilde{\lambda}$ by λ and we only have n terms rather than $n + 1$ terms. This gives

$$\mathbb{E}[A_{t-1} - A_t] = \eta\lambda(A_{t-1} + \sum_{i=1}^n \frac{1}{q_i} (-\|u_i - \hat{a}_i\|_2^2 + (1 - \beta_i)\|m_i\|_2^2)).$$

Note that $g^*(v)$ is 1-smooth, follow the same analysis of Step 2 in the proof of Theorem 1, specifically the derivation in (21), we obtain

$$\frac{1}{2}(B_{t-1} - B_t) \geq \langle w^{t-1} - \hat{w}, v^{t-1} - v^t \rangle - \frac{1}{2}\|v^{t-1} - v^t\|_2^2.$$

Combining the above two equations, we have

$$\begin{aligned} \mathbb{E}(C_{t-1} - C_t) &\geq c_a\eta\lambda A_{t-1} - c_a\eta\lambda \sum_{i=1}^n \frac{1}{q_i} \|u_i - \hat{a}_i\|_2^2 + 2c_b\eta \langle w^{t-1} - \hat{w}, \nabla f(w^{t-1}) + \lambda v^{t-1} \rangle \\ &\quad + \sum_{i=1}^n \frac{1}{q_i} \|m_i\|_2^2 (c_a\eta\lambda(1 - \beta_i) - \frac{c_b\eta^2}{n^2}). \end{aligned} \tag{31}$$

We can choose $\eta \leq \frac{q_i}{2\lambda}$ and $\frac{c_b}{c_a} = \frac{\lambda n^2}{2\eta}$ so that $\beta \leq 1/2$, and the term $\sum_{i=1}^n \frac{1}{q_i} \|m_i\|_2^2 (c_a\eta\lambda(1 - \beta_i) - \frac{c_b\eta^2}{n^2})$ is non-negative. According to the definition of q_i , we know $q_i \geq \frac{1}{2n}$, thus we can choose $\eta \leq \frac{1}{4n}$ to ensure that $\eta \leq \frac{q_i}{2\lambda}$.

We can show that

$$\begin{aligned} &\langle w^{t-1} - \hat{w}, \nabla f(w^{t-1}) + \lambda v^{t-1} \rangle - \frac{\lambda}{2} B_{t-1} - (F(w^{t-1}) - F(\hat{w})) \\ &= f(\hat{w}) - f(w^{t-1}) - \langle \hat{w} - w^{t-1}, \nabla f(w^{t-1}) \rangle \geq 0 \end{aligned} \tag{32}$$

This follows from the same derivation as in Equation (24), modulus replacing $\tilde{\lambda}$, $\tilde{g}(w)$ and $\phi_i(w)$ by λ , $g(w)$ and $f_i(w)$ correspondingly. The last step follows immediately from the convexity of $f(w)$.

The proof now proceed with discussing two cases separately: (1). each $f_i(w)$ is convex. And (2). $f_i(w)$ is not necessarily convex, but the sum $f(w)$ is convex.

Case 1: Since each $f_i(w)$ is convex in this case, to bound $\sum_{i=1}^n \frac{1}{q_i} \|u_i - \hat{a}_i\|_2^2$ is easier, via the following.

$$\begin{aligned} \sum_{i=1}^n \frac{1}{q_i} \|u_i - \hat{a}_i\|_2^2 &= \sum_{i=1}^n \frac{1}{q_i} \|\nabla f_i(w^{t-1}) - \nabla f_i(\hat{w})\|_2^2 \\ &\leq (2 \max_i \frac{L_i}{q_i}) \sum_{i=1}^n (f_i(w^{t-1}) - f(\hat{w}) - \langle \nabla f_i(\hat{w}), w^{t-1} - \hat{w} \rangle) \\ &\leq (2 \max_i \frac{L_i}{q_i}) n(F(w^{t-1}) - F(\hat{w})), \end{aligned} \quad (33)$$

where the first inequality follows from Lemma 1, and the second one follows from convexity of $g(w)$ and optimality condition of \hat{w} .

The definition of q_i in the Algorithm 1 implies that for every i ,

$$\frac{L_i}{q_i} = 2n\bar{L} \frac{L_i}{L_i + L} \leq 2n\bar{L}.$$

Substitute this into the corresponding terms in (31), we have

$$\begin{aligned} &\mathbb{E}(C_{t-1} - C_t) \\ &\geq c_a \eta \lambda A_{t-1} + c_b \eta \lambda B_{t-1} + 2c_b \eta (F(w^{t-1}) - F(\hat{w})) - 4c_a \eta \lambda n^2 \bar{L} (F(w^{t-1}) - F(\hat{w})). \end{aligned} \quad (34)$$

Remind $c_b = \frac{c_a \lambda n^2}{2\eta}$, using the choice $\eta \leq \frac{1}{4L}$, we have

$$E(C_t) \leq (1 - \eta \lambda) C_{t-1},$$

where $\eta = \min\{\frac{1}{4L}, \frac{1}{4\lambda n}\}$.

Case 2: Using strong convexity of $F(\cdot)$, we have

$$\begin{aligned} c_a \eta \lambda \sum_{i=1}^n \frac{1}{q_i} \|u_i - \hat{a}_i\|_2^2 &= c_a \eta \lambda \sum_{i=1}^n \frac{1}{q_i} \|\nabla f_i(w^{t-1}) - \nabla f_i(\hat{w})\|_2^2 \\ &\leq c_a \eta \lambda \sum_{i=1}^n \frac{L_i^2}{q_i} \|w^{t-1} - \hat{w}\|_2^2 \leq 2c_a \eta (F(w^{t-1}) - F(\hat{w})) \sum_{i=1}^n \frac{L_i^2}{q_i}, \end{aligned} \quad (35)$$

where the first inequality uses the smoothness of $f_i(\cdot)$, and the second one holds from the strong convexity of $F(\cdot)$ and optimal condition of \hat{w} .

Using Equation (32), we have

$$\langle w^{t-1} - \hat{w}, \nabla f(w^{t-1}) + \lambda v^{t-1} \rangle \geq \frac{\lambda}{2} B_{t-1} + (F(w^{t-1}) - F(\hat{w})).$$

Then replace the corresponding terms in (31) we have

$$\mathbb{E}(C_{t-1} - C_t) \geq \eta \lambda C_{t-1} + 2\eta (c_b - c_a \sum_{i=1}^n \frac{L_i^2}{q_i}) (F(w^{t-1}) - F(\hat{w})).$$

Since we know $\frac{L_i}{q_i} \leq 2n\bar{L}$, we have

$$\mathbb{E}(C_{t-1} - C_t) \geq \eta \lambda C_{t-1} + 2\eta (c_b - 2n^2 \bar{L}^2 c_a) (F(w^{t-1}) - F(\hat{w})).$$

The last term is no-negative if $\frac{c_b}{c_a} \geq 2n^2 \bar{L}^2$. Since we choose $\frac{c_b}{c_a} = \frac{\lambda n^2}{2\eta}$, this condition is satisfied as we choose $\eta \leq \frac{\lambda}{4L^2}$. \square

A.3 Proof on non-convex $F(w)$

The proof of the non-convex case follows a similar line as that of the convex case. To avoid redundancy, we focus on pointing out the difference in the proof. We start with some technical lemmas. The following lemma is adapted from Lemma 6 of Loh and Wainwright [2013], which we present for completeness.

Lemma 6. *For any vector $w \in \mathbb{R}^p$, let A denote the index set of its s largest elements in magnitude, under assumption of $d_{\lambda,\mu}$ in Section 2.3 in the main body of paper, we have*

$$d_{\lambda,\mu}(w_A) - d_{\lambda,\mu}(w_{A^c}) \leq \lambda L_d (\|w_A\|_1 - \|w_{A^c}\|_1).$$

Moreover, for an arbitrary vector $w \in \mathbb{R}^p$, we have

$$d_{\lambda,\mu}(w^*) - d_{\lambda,\mu}(w) \leq \lambda L_d (\|v_A\|_1 - \|v_{A^c}\|_1),$$

where $v = w - w^*$ and w^* is s sparse.

The next lemma is non-convex counterparts of Lemma 3 and Lemma 4.

Lemma 7. *Suppose $d_{\lambda,\mu}(\cdot)$ satisfies the assumptions in Section 2.3 in the main body of paper, w^* is feasible, $\lambda L_d \geq 8\rho\theta\frac{\log p}{n}$, $\lambda \geq \frac{4}{L_d}\|\nabla f(w^*)\|_\infty$, and there exists ξ, T such that*

$$F(w^t) - F(\hat{w}) \leq \xi, \forall t > T.$$

Then for any $t > T$, we have

$$\|w^t - \hat{w}\|_1 \leq 4\sqrt{s}\|w^t - \hat{w}\|_2 + 8\sqrt{s}\|w^* - \hat{w}\|_2 + 2\min\left(\frac{\xi}{\lambda L_d}, \rho\right).$$

Proof. For an arbitrary feasible w , define $\Delta = w - w^*$. Suppose we have $F(w) - F(\hat{w}) \leq \xi$, since we know $F(\hat{w}) \leq F(w^*)$ so we have $F(w) \leq F(w^*) + \xi$, which implies

$$f(w^* + \Delta) + d_{\lambda,\mu}(w^* + \Delta) \leq f(w^*) + d_{\lambda,\mu}(w^*) + \xi.$$

Subtract $\langle \nabla f(w^*), \Delta \rangle$ and use the RSC condition (Recall we have $\tau = \theta\frac{\log p}{n}$ in the assumption of Theorem 2) we have

$$\begin{aligned} & \frac{\kappa}{2}\|\Delta\|_2^2 - \theta\frac{\log p}{n}\|\Delta\|_1^2 + d_{\lambda,\mu}(w^* + \Delta) - d_{\lambda,\mu}(w^*) \\ & \leq \xi - \langle \nabla f(w^*), \Delta \rangle \\ & \leq \xi + \|\nabla f(w^*)\|_\infty \|\Delta\|_1, \end{aligned} \tag{36}$$

where the last inequality holds from Holder's inequality.

Rearranging terms and use the fact that $\|\Delta\|_1 \leq 2\rho$ (by feasibility of w and w^*) and we assume that $\lambda L_d \geq 8\rho\theta\frac{\log p}{n}$, $\lambda \geq \frac{4}{L_d}\|\nabla f(w^*)\|_\infty$, we get

$$\xi + \frac{1}{2}\lambda L_d \|\Delta\|_1 + d_{\lambda,\mu}(w^*) - d_{\lambda,\mu}(w^* + \Delta) \geq \frac{\kappa}{2}\|\Delta\|_2^2 \geq 0.$$

By Lemma 6, we have

$$d_{\lambda,\mu}(w^*) - d_{\lambda,\mu}(w) \leq \lambda L_d (\|\Delta_A\|_1 - \|\Delta_{A^c}\|_1),$$

where A is the set of indices of the top s components of Δ in magnitude, which thus leads to

$$\frac{3\lambda L_d}{2}\|\Delta_A\|_1 - \frac{\lambda L_d}{2}\|\Delta_{A^c}\|_1 + \xi \geq 0.$$

Consequently

$$\|\Delta\|_1 \leq \|\Delta_A\|_1 + \|\Delta_{A^c}\|_1 \leq 4\|\Delta_A\|_1 + \frac{2\xi}{\lambda L_d} \leq 4\sqrt{s}\|\Delta\|_2 + \frac{2\xi}{\lambda L_d}.$$

Combining with the fact $\|\Delta\|_1 \leq 2\rho$, we obtain

$$\|\Delta\|_1 \leq 4\sqrt{s}\|\Delta\|_2 + 2\min\left\{\frac{\xi}{\lambda L_d}, \rho\right\}.$$

So we have $\|w^t - w^*\|_1 \leq 4\sqrt{s}\|w^t - w^*\|_2 + 2\min\left\{\frac{\xi}{\lambda L_d}, \rho\right\}$.

Notice $F(w^*) - F(\hat{w}) \geq 0$, so following same steps and set $\xi = 0$ we have $\|\hat{w} - w^*\|_1 \leq 4\sqrt{s}\|\hat{w} - w^*\|_2$.

Combining the two together, we get

$$\|w^t - \hat{w}\|_1 \leq \|w^t - w^*\|_1 + \|w^* - \hat{w}\|_1 \leq 4\sqrt{s}\|w^t - \hat{w}\|_2 + 8\sqrt{s}\|w^* - \hat{w}\|_2 + 2\min\left(\frac{\xi}{\lambda L_d}, \rho\right).$$

□

Now we provide a counterpart of Lemma 5 in the non-convex case.

Lemma 8. *Under the same assumptions as those of Lemma 7, we have*

$$F(w^t) - F(\hat{w}) \geq \frac{\tilde{\kappa}}{2}\|w^t - \hat{w}\|_2^2 - \epsilon^2(\Delta^*, s);$$

$$\text{and } \phi(\hat{w}) - \phi(w^t) - \langle \nabla \phi(w^t), \hat{w} - w^t \rangle \geq \left[\frac{\tilde{\kappa} - \tilde{\lambda}}{2}\|w^t - \hat{w}\|_2^2 - \epsilon^2(\Delta^*, s)\right],$$

where $\tilde{\kappa} = \kappa - \mu - 64s\tau$, $\Delta^* = \hat{w} - w^*$, and $\epsilon^2(\Delta^*, s) = 2\tau(8\sqrt{s}\|\hat{w} - w^*\|_2 + 2\min(\frac{\xi}{\lambda L_d}, \rho))^2$.

Proof. Notice that

$$\begin{aligned} & F(w^t) - F(\hat{w}) \\ &= f(w^t) - f(\hat{w}) - \frac{\mu}{2}\|w^t\|_2^2 + \frac{\mu}{2}\|\hat{w}\|_2^2 + \lambda d_\lambda(w^t) - \lambda d_\lambda(\hat{w}) \\ &\geq \langle \nabla f(\hat{w}), w^t - \hat{w} \rangle + \frac{\kappa}{2}\|w^t - \hat{w}\|_2^2 - \langle \mu \hat{w}, w^t - \hat{w} \rangle - \frac{\mu}{2}\|w^t - \hat{w}\|_2^2 \\ &\quad + \lambda d_\lambda(w^t) - \lambda d_\lambda(\hat{w}) - \tau\|w^t - \hat{w}\|_1^2 \\ &\geq \langle \nabla f(\hat{w}), w^t - \hat{w} \rangle + \frac{\kappa}{2}\|w^t - \hat{w}\|_2^2 - \langle \mu \hat{w}, w^t - \hat{w} \rangle - \frac{\mu}{2}\|w^t - \hat{w}\|_2^2 \\ &\quad + \lambda \langle \partial d_\lambda(\hat{w}), w^t - \hat{w} \rangle - \tau\|w^t - \hat{w}\|_1^2 \\ &= \frac{\kappa - \mu}{2}\|w^t - \hat{w}\|_2^2 - \tau\|w^t - \hat{w}\|_1^2, \end{aligned} \tag{37}$$

where the first inequality uses RSC condition, the second inequality uses the convexity of $d_\lambda(w)$, and the last equality holds from the optimality condition of \hat{w} .

Using Lemma 7, and the inequality $(a + b)^2 \leq 2a^2 + 2b^2$, we have

$$\begin{aligned} \|w^t - \hat{w}\|_1^2 &\leq (4\sqrt{s}\|w^t - \hat{w}\|_2 + 8\sqrt{s}\|w^* - \hat{w}\|_2 + 2\min(\frac{\xi}{\lambda L_d}, \rho))^2 \\ &\leq 32s\|w^t - \hat{w}\|_2^2 + 2(8\sqrt{s}\|\hat{w} - w^*\|_2 + 2\min(\frac{\xi}{\lambda L_d}, \rho))^2. \end{aligned} \tag{38}$$

Substitute this into Equation (37), we obtain

$$F(w^t) - F(\hat{w}) \geq \left(\frac{\kappa - \mu}{2} - 32s\tau\right)\|w^t - \hat{w}\|_2^2 - 2\tau(8\sqrt{s}\|\hat{w} - w^*\|_2 + 2\min(\frac{\xi}{\lambda L_d}, \rho))^2.$$

Recall

$$\phi(w) = f(w) - \frac{\tilde{\lambda}}{2}\|w\|_2^2 - \frac{\mu}{2}\|w\|_2^2.$$

So we have

$$\begin{aligned} & \phi(\hat{w}) - \phi(w^t) - \langle \hat{w} - w^t, \nabla \phi(w^t) \rangle \\ &= f(\hat{w}) - f(w^t) - \langle \nabla f(w^t), \hat{w} - w^t \rangle - \frac{\tilde{\lambda} + \mu}{2}\|\hat{w} - w^t\|_2^2 \\ &\geq \frac{\kappa}{2}\|\hat{w} - w^t\|_2^2 - \tau\|\hat{w} - w^t\|_1^2 - \frac{\tilde{\lambda} + \mu}{2}\|\hat{w} - w^t\|_2^2. \end{aligned} \tag{39}$$

Then use the upper bound of $\|w^t - \hat{w}\|_1^2$ in (38) and rearrange terms we establish the lemma. \square

We are now ready to prove the main theorem of the non-convex case, i.e., Theorem 2.

Proof of Theorem 2. Remind in the non-convex case, for $i = 1, \dots, n$, the definition of $\phi_i(x)$ is the same as in the convex case. The difference is for $\phi_{n+1}(w)$, where $\phi_{n+1}(w)$ is defined as follows:

$$\phi_{n+1}(w) = -\frac{(n+1)(\tilde{\lambda} + \mu)}{2}\|w\|_2^2.$$

Recall the definition of $\tilde{g}(w)$ is as follows

$$\tilde{g}(w) = \frac{1}{2}\|w\|_2^2 + \frac{\lambda}{\tilde{\lambda}}d_\lambda(w).$$

Following similar steps as in the proof of Theorem 1, we have

$$\begin{aligned} E(C_{t-1} - C_t) &\geq c_a\eta\tilde{\lambda}A_{t-1} - c_a\eta\tilde{\lambda}\sum_{i=1}^{n+1}\frac{1}{q_i}\|u_i - \hat{a}_i\|_2^2 + c_b\eta\tilde{\lambda}B_{t-1} + 2c_b\eta(F(w^{t-1}) - F(\hat{w})) \\ &+ \sum_{i=1}^{n+1}\frac{1}{q_i}\|m_i\|_2^2(c_a\eta\tilde{\lambda}(1 - \beta_i) - \frac{c_b\eta^2}{(n+1)^2}) + 2c_b\eta[\phi(\hat{w}) - \phi(w^{t-1}) - \langle \hat{w} - w^{t-1}, \nabla \phi(w^{t-1}) \rangle]. \end{aligned} \tag{40}$$

Following the same steps as in Equation (27), and replace $\tilde{\lambda}$ by $\tilde{\lambda} + \mu$ we have

$$\frac{1}{q_{n+1}}\|\nabla \phi_{n+1}(w^{t-1}) - \nabla \phi_{n+1}(\hat{w})\|_2^2 \leq 2(n+1)\frac{\tilde{L}_{n+1}}{q_{n+1}}\frac{\tilde{\lambda} + \mu}{2}\|w^{t-1} - \hat{w}\|_2^2. \tag{41}$$

Similar as the convex case we then divide the time step $t = 1, 2, \dots$ into several epochs, i.e., $([T_0, T_1], [T_1, T_2], \dots)$. At the end of each epoch j , we prove that C_t decreases with a linear rate until the optimality gap $F(w^t) - F(\hat{w})$ decrease to some tolerance ξ_j .

We then apply Lemma 8 and using the assumption that $\tilde{\kappa} \geq \tilde{\lambda} + \mu$ to relate $\|w^{t-1} - \hat{w}\|_2^2$ to $F(w^t) - F(\hat{w})$:

$$\frac{1}{q_{n+1}}\|\nabla \phi_{n+1}(w^{t-1}) - \nabla \phi_{n+1}(\hat{w})\|_2^2 \leq 2(n+1)\frac{\tilde{L}_{n+1}}{q_{n+1}}(F(w^{t-1}) - F(\hat{w}) + \epsilon_j^2(\Delta^*, s)).$$

Thus

$$\sum_{i=1}^{n+1} \frac{1}{q_i} \|u_i - \hat{a}_i\|_2^2 \leq 8(n+1)^2 \tilde{L}(F(w^{t-1}) - F(\hat{w})) + 4(n+1)^2 \tilde{L}\epsilon_j^2(\Delta^*, s). \quad (42)$$

We choose $2c_b = 16c_a\tilde{\lambda}(n+1)^2\tilde{L}$, and use the fact that $\frac{c_b}{c_a} = \frac{\tilde{\lambda}(n+1)^2}{2\eta}$, we have $\eta = \frac{1}{16\tilde{L}}$. Combine all pieces together we get

$$\begin{aligned} \mathbb{E}(C_{t-1} - C_t) &\geq c_a\eta\tilde{\lambda}A_{t-1} - 8c_a\eta\tilde{\lambda}(n+1)^2\tilde{L}(F(w^{t-1}) - F(\hat{w})) - 4c_a\eta\tilde{\lambda}(n+1)^2\tilde{L}\epsilon_j^2(\Delta^*, s) \\ &\quad + c_b\eta\tilde{\lambda}B_{t-1} + 2c_b\eta(F(w^{t-1}) - F(\hat{w})) + 2c_b\eta(\phi(\hat{w}) - \phi(w^{t-1}) - \langle \hat{w} - w^{t-1}, \nabla\phi(w^{t-1}) \rangle) \\ &\geq \eta\tilde{\lambda}C_{t-1} + \eta(2c_b - 8c_a\tilde{\lambda}(n+1)^2\tilde{L})(F(w^{t-1}) - F(\hat{w})) - 4c_a\eta\tilde{\lambda}(n+1)^2\tilde{L}\epsilon_j^2(\Delta^*, s) \\ &\quad + (c_b\eta(\tilde{\kappa} - \tilde{\lambda})\|w^t - \hat{w}\|_2^2 - 2c_b\eta\epsilon_j^2(\Delta^*, s)) \\ &\geq \eta\tilde{\lambda}C_{t-1} + \eta c_b(F(w^{t-1}) - F(\hat{w})) - 3c_b\eta\epsilon_j^2(\Delta^*, s), \end{aligned} \quad (43)$$

where the second inequality uses Lemma 8 and the assume $\tilde{\kappa} \geq \tilde{\lambda}$.

The rest of the proofs are almost identical to the convex one. In the first epoch, we have

$$\epsilon_1^2(\Delta^*, s) = 2\tau(\delta_{stat} + 2\rho)^2, \xi_1 = 6\tau(\delta_{stat} + 2\rho)^2$$

Then we choose T_1 such that

$$F(w^{T_1-1}) - F(\hat{w}) \geq 3\epsilon_1^2(\Delta^*, s) \quad \text{and} \quad F(w^{T_1}) - F(\hat{w}) \leq 3\epsilon_1^2(\Delta^*, s).$$

In the second epoch we can use the same argument.

$$\epsilon_2^2(\Delta^*, s) = 2\tau_\sigma(\delta_{stat} + \delta_1)^2 \quad \text{where} \quad \delta_1 = 2\frac{\xi_1}{\lambda L_d}.$$

Repeat this strategy on every epoch. The only difference from the proof of the convex case is that we now replace λ by λL_d .

We use the same argument to prove ξ_k is a decreasing sequence and conclude the proof. \square

A.4 Proof of Corollaries

We now prove the corollaries that instantiate our main theorems to different statistical estimators.

Proof of corollary on Lasso, i.e., corollary 1. To begin with, we present the following lemma of the RSC proved in Raskutti et al. [2010] and then use it in the case of Lasso.

Lemma 9. *If each data point x_i is i.i.d randomly sampled from the distribution $N(0, \Sigma)$, then there exist universal constants c_0 and c_1 such that*

$$\frac{\|X\Delta\|_2^2}{n} \geq \frac{1}{2}\|\Sigma^{1/2}\Delta\|_2^2 - c_1\nu(\Sigma)\frac{\log p}{n}\|\Delta\|_1^2, \quad \text{for all } \Delta \in \mathbb{R}^p,$$

with probability at least $1 - \exp(-c_0n)$. Here X is the data matrix where each row is data point x_i .

Since w^* is supported on a subset S with cardinality s , we choose

$$\mathcal{B}(S) = \{w \in \mathbb{R}^p | w_j = 0 \text{ for all } j \notin S\}.$$

It is straightforward to choose $\mathcal{A}(S) = \mathcal{B}(S)$ and notice that $w^* \in \mathcal{A}(S)$.

In Lasso, $f(w) = \frac{1}{2n} \|y - Xw\|_2^2$, and it is easy to verify that

$$f(w + \Delta) - f(w) - \langle \nabla f(w), \Delta \rangle \geq \frac{1}{2n} \|X\Delta\|_2^2 \geq \frac{1}{4} \|\Sigma^{1/2} \Delta\|_2^2 - \frac{c_1}{2} \nu(\Sigma) \frac{\log p}{n} \|\Delta\|_1^2.$$

Notice that $g(\cdot)$ is $\|\cdot\|_1$ in Lasso, thus $\Psi(\mathcal{B}) = \sup_{w \in \mathcal{B} \setminus \{0\}} \frac{\|w\|_1}{\|w\|_2} = \sqrt{s}$. So we have

$$\tilde{\kappa} = \frac{1}{2} \sigma_{\min}(\Sigma) - 64c_1 \nu(\Sigma) \frac{s \log p}{n}.$$

On the other hand, the tolerance is

$$\begin{aligned} \delta &= 24\tau(8\Psi(\mathcal{B})\|\hat{w} - w^*\|_2 + 8g(w_{\mathcal{A}^\perp}^*))^2 \\ &= c_2 \nu(\Sigma) \frac{s \log p}{n} \|\hat{w} - w^*\|_2^2, \end{aligned} \tag{44}$$

where we use the fact that $w^* \in \mathcal{A}(S)$ which implies $g(w_{\mathcal{A}^\perp}^*) = 0$.

The last piece to check is that $\lambda \geq 2g^*(\nabla f(w^*))$. In Lasso we have $g^*(\cdot) = \|\cdot\|_\infty$. Using the fact that $y_i = x_i^T w^* + \xi_i$, this is equivalent to require $\lambda \geq \frac{2}{n} \|X^T \xi\|_\infty$. Thus, by our choice of λ that $\lambda \geq 6\sigma \sqrt{\frac{\log p}{n}}$, the condition is satisfied invoking the following inequality which holds by applying the tail bound on the Gaussian variable and the union bound to get

$$\mathbb{P} \left(\frac{2}{n} \|X^T \xi\|_\infty \leq 6\sigma \sqrt{\frac{\log p}{n}} \right) \geq 1 - \exp(-3 \log p).$$

□

Proof of corollary on Group Lasso, i.e., corollary 2. We use the following fact on the RSC condition of the Group Lasso Negahban et al. [2009]Negahban et al. [2012]. As each data point x_i is i.i.d random sampled from the distribution $N(0, \Sigma)$, then there exists strictly positive constant (σ_1, σ_2) which just depends on Σ such that

$$\frac{\|X\Delta\|_2^2}{2n} \geq \sigma_1(\Sigma) \|\Delta\|_2^2 - \sigma_2(\Sigma) \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{3 \log N_{\mathcal{G}}}{n}} \right)^2 \|\Delta\|_{\mathcal{G},2}^2, \quad \text{for all } \Delta \in \mathbb{R}^p$$

with probability at least $1 - c_3 \exp(-c_4 n)$.

Recall we define the subspace

$$\mathcal{A}(S_{\mathcal{G}}) = \{w | w_{G_i} = 0 \text{ for all } i \notin S_{\mathcal{G}}\},$$

and $\mathcal{A}(S_{\mathcal{G}}) = \mathcal{B}(S_{\mathcal{G}})$, where $S_{\mathcal{G}}$ corresponds to non-zero group of w^* . The subspace compatibility can be computed by

$$\Psi(\mathcal{B}) = \sup_{w \in \mathcal{B} \setminus \{0\}} \frac{\|w\|_{\mathcal{G},2}}{\|w\|_2} = \sqrt{s_{\mathcal{G}}}.$$

The effective RSC parameter is given by

$$\tilde{\kappa} = \sigma_1(\Sigma) - 64\sigma_2(\Sigma)s_{\mathcal{G}} \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{3 \log N_{\mathcal{G}}}{n}} \right)^2.$$

We then check the requirement on λ holds. Since Group lasso is $\ell_{1,2}$ grouped norm, its dual norm of it is $(\infty, 2)$ grouped norm. So we need

$$\lambda \geq 2 \max_{i=1, \dots, N_{\mathcal{G}}} \left\| \frac{1}{n} (X^T \xi)_{G_i} \right\|_2.$$

Using Lemma 5 in Negahban et al. [2009], we know

$$\max_{i=1, \dots, N_{\mathcal{G}}} \left\| \frac{1}{n} (X^T \xi)_{G_i} \right\|_2 \leq 2\sigma \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{\log N_{\mathcal{G}}}{n}} \right)$$

with probability at least $1 - 2 \exp(-2 \log N_{\mathcal{G}})$. Thus it suffices to choose $\lambda \geq 4\sigma(\sqrt{\frac{m}{n}} + \sqrt{\frac{\log N_{\mathcal{G}}}{n}})$, as suggested in the corollary.

Finally, the tolerance can be computed as

$$\delta = c_3 s_{\mathcal{G}} \sigma_2(\Sigma) \left(\sqrt{\frac{m}{n}} + \sqrt{\frac{3 \log N_{\mathcal{G}}}{n}} \right)^2 \|\hat{w} - w^*\|_2^2.$$

□

Proof of corollary on Linear regression with SCAD regularizer, i.e., corollary 3. The proof is similar to that of Lasso. Recall in the proof of the Lasso example we have

$$\|\nabla f(w^*)\|_{\infty} = \frac{1}{n} \|X^T \xi\|_{\infty} \leq 3\sigma \sqrt{\frac{\log p}{n}},$$

and the RSC condition is

$$\frac{\|X\Delta\|_2^2}{n} \geq \frac{1}{2} \|\Sigma^{1/2} \Delta\|_2^2 - c_1 \nu(\Sigma) \frac{\log p}{n} \|\Delta\|_1^2.$$

Recall $\mu = \frac{1}{\zeta-1}$ and $L_d = 1$, we establish the result. □

Proof of corollary on corrected Lasso, i.e., corollary 4. Notice

$$\|\nabla f(w^*)\|_{\infty} = \|\hat{\Gamma} w^* - \hat{\gamma}\|_{\infty} = \|\hat{\gamma} - \Sigma w^* + (\Sigma - \hat{\Gamma}) w^*\|_{\infty} \leq \|\hat{\gamma} - \Sigma w^*\|_{\infty} + \|(\Sigma - \hat{\Gamma}) w^*\|_{\infty}.$$

As shown in Lemma 2 of Loh and Wainwright [2011], both terms are bounded by $c_1 \varphi \sqrt{\frac{\log p}{n}}$ with probability at least $1 - c_1 \exp(-c_2 \log p)$, where $\varphi = (\sqrt{\sigma_{\max}(\Sigma)} + \sqrt{\gamma_{\zeta}})(\sigma + \sqrt{\gamma_{\zeta}} \|w^*\|_2)$.

To obtain the RSC condition, we apply Lemma 1 in Loh and Wainwright [2011], to get

$$\frac{1}{n} \Delta^T \hat{\Gamma} \Delta \geq \frac{\sigma_{\min}(\Sigma)}{2} \|\Delta\|_2^2 - c_3 \sigma_{\min}(\Sigma) \max \left(\left(\frac{\sigma_{\max}(\Sigma) + \gamma_{\zeta}}{\sigma_{\min}(\Sigma)} \right)^2, 1 \right) \frac{\log p}{n} \|\Delta\|_1^2,$$

with probability at least $1 - c_4 \exp \left(-c_5 n \min \left(\frac{\sigma_{\min}^2(\Sigma)}{(\sigma_{\max}(\Sigma) + \gamma_{\zeta})^2}, 1 \right) \right)$.

Combine these pieces together, we establish the result. □

References

- Alekh Agarwal, Sahand Negahban, and Martin J Wainwright. Fast global convergence rates of gradient methods for high-dimensional statistical recovery. In *Advances in Neural Information Processing Systems*, pages 37–45, 2010.
- Zeyuan Allen-Zhu and Yang Yuan. Even faster accelerated coordinate descent using non-uniform sampling. *arXiv preprint arXiv:1512.09103*, 2015a.
- Zeyuan Allen-Zhu and Yang Yuan. Improved svrg for non-strongly-convex or sum-of-non-convex objectives. *arXiv preprint arXiv:1506.01972*, 2015b.
- Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT’2010*, pages 177–186. Springer, 2010.
- Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in Neural Information Processing Systems*, pages 1646–1654, 2014.
- John Duchi and Yoram Singer. Efficient online and batch learning using forward backward splitting. *Journal of Machine Learning Research*, 10(Dec):2899–2934, 2009.
- Celestine Dünnér, Simone Forte, Martin Takáč, and Martin Jaggi. Primal-dual rates and certificates. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning-Volume 48*, pages 783–792. JMLR. org, 2016.
- Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- I. Guyon. A phamacology dataset, 06 2008. URL <http://www.causality.inf.ethz.ch/data/SID0.html>.
- D. Harrison and D.L. Rubinfeld. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- Martin Jaggi, Virginia Smith, Martin Takáč, Jonathan Terhorst, Sanjay Krishnan, Thomas Hofmann, and Michael I Jordan. Communication-efficient distributed dual coordinate ascent. In *Advances in Neural Information Processing Systems*, pages 3068–3076, 2014.
- David D Lewis, Yiming Yang, Tony G Rose, and Fan Li. Rcv1: A new benchmark collection for text categorization research. *Journal of machine learning research*, 5 (Apr):361–397, 2004.
- Xingguo Li, Tuo Zhao, Raman Arora, Han Liu, and Jarvis Haupt. Stochastic variance reduced optimization for nonconvex sparse learning. *arXiv preprint arXiv:1605.02711*, 2016.
- Ji Liu and Stephen J Wright. Asynchronous stochastic coordinate descent: Parallelism and convergence properties. *SIAM Journal on Optimization*, 25(1):351–376, 2015.

- Po-Ling Loh and Martin J Wainwright. High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. In *Advances in Neural Information Processing Systems*, pages 2726–2734, 2011.
- Po-Ling Loh and Martin J Wainwright. Regularized m -estimators with nonconvexity: Statistical and algorithmic theory for local optima. In *Advances in Neural Information Processing Systems*, pages 476–484, 2013.
- Dhruv Mahajan, S Sathiya Keerthi, and S Sundararajan. A distributed block coordinate descent method for training l_1 regularized linear classifiers. *arXiv preprint arXiv:1405.4544*, 2014.
- S Negahban, P Ravikumar, MJ Wainwright, and B Yu. Supplement to a unified framework for high-dimensional analysis of m -estimators with decomposable regularizers., 2012.
- Sahand Negahban, Bin Yu, Martin J Wainwright, and Pradeep K Ravikumar. A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. In *Advances in Neural Information Processing Systems*, pages 1348–1356, 2009.
- Yu Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.
- Yu Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.
- Danil Prokhorov. Ijcnn 2001 neural network competition, 2001.
- Chao Qu, Yan Li, and Huan Xu. Linear convergence of svrg in statistical estimation. *arXiv preprint arXiv:1611.01957*, 2016.
- Chao Qu, Yan Li, and Huan Xu. Saga and restricted strong convexity. *arXiv preprint arXiv:1702.05683*, 2017.
- Zheng Qu and Peter Richtárik. Coordinate descent with arbitrary sampling i: Algorithms and complexity. *arXiv preprint arXiv:1412.8060*, 2014.
- Zheng Qu and Peter Richtárik. Stochastic dual coordinate ascent with adaptive probabilities. 2015.
- Zheng Qu, Peter Richtárik, and Tong Zhang. Quartz: Randomized dual coordinate ascent with arbitrary sampling. In *Advances in Neural Information Processing Systems*, pages 865–873, 2015.
- Ariadna Quattoni, Xavier Carreras, Michael Collins, and Trevor Darrell. An efficient projection for l_1 , infinity regularization. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 857–864. ACM, 2009.
- Garvesh Raskutti, Martin J Wainwright, and Bin Yu. Restricted eigenvalue properties for correlated gaussian designs. *Journal of Machine Learning Research*, 11(Aug):2241–2259, 2010.

- Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic average gradient. *arXiv preprint arXiv:1309.2388*, 2013.
- Shai Shalev-Shwartz. Sdca without duality, regularization, and individual convexity. *arXiv preprint arXiv:1602.01582*, 2016.
- Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14(Feb):567–599, 2013.
- Shai Shalev-Shwartz and Tong Zhang. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. In *ICML*, pages 64–72, 2014.
- Grzegorz Swirszcz, Naoki Abe, and Aurelie C Lozano. Grouped orthogonal matching pursuit for variable selection and prediction. In *Advances in Neural Information Processing Systems*, pages 1150–1158, 2009.
- Berwin A Turlach, William N Venables, and Stephen J Wright. Simultaneous variable selection. *Technometrics*, 47(3):349–363, 2005.
- Sara A Van De Geer, Peter Bühlmann, et al. On the conditions used to prove oracle results for the lasso. *Electronic Journal of Statistics*, 3:1360–1392, 2009.
- A. Vedaldi and B. Fulkerson. VLFeat: An open and portable library of computer vision algorithms. <http://www.vlfeat.org/>, 2008.
- Stephen J Wright. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34, 2015.
- Shuo Xiang, Tao Yang, and Jieping Ye. Simultaneous feature and feature group selection through hard thresholding. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 532–541. ACM, 2014.
- Lin Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11(Oct):2543–2596, 2010.
- Lin Xiao and Tong Zhang. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.
- Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.
- Cun-Hui Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of statistics*, pages 894–942, 2010.
- Tuo Zhao, Mo Yu, Yiming Wang, Raman Arora, and Han Liu. Accelerated mini-batch randomized block coordinate descent method. In *Advances in neural information processing systems*, pages 3329–3337, 2014.